

12

20000801207

TR-283

STUDIES IN MATHEMATICAL MODELS OF HUMAN
DECISIONMAKING IN GAMING SITUATIONS


Reproduced From
Best Available Copy

AD-A166 698

DTIC FILE COPY

DTIC
ELECTE
APR 16 1988
S D D

This document has been approved
for public release and sale; its
distribution is unlimited.



ALPHATECH
INC.
2 BURLINGTON EXECUTIVE CENTER
111 MIDDLESEX TURNPIKE
BURLINGTON, MA 01803
617-273-3388

86 4 16 007

(2)

DTIC
ELECTE
APR 16 1986
S D

TR-283

**STUDIES IN MATHEMATICAL MODELS OF HUMAN
DECISIONMAKING IN GAMING SITUATIONS**

By

Dr. David A. Castanon

April 1986

Prepared for:

**Mr. J. Randolph Simpson
Code 411
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217**

Contract Number: N00014-84-C-0458

**ALPHATECH, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, MA 01803
(617) 273-3388**

**Document has been approved
for public release and sale; its
distribution is unlimited.**

CONTENTS

	<u>Page</u>
EXECUTIVE SUMMARY.	1
APPENDIX A - AN OVERVIEW OF RESULTS IN HUMAN DECISIONMAKING.	A-0
APPENDIX B - TECHNICAL REPORT TR-172: ASYMPTOTIC AGREEMENT IN DISTRIBUTED ESTIMATION WITH INCONSISTENT BELIEFS.	B-0
APPENDIX C - CONSENSUS IN DISTRIBUTED ESTIMATION	C-0
APPENDIX D - FURTHER RESULTS ON THE CONSENSUS PROBLEM.	D-0
APPENDIX E - INFORMATIONAL ASPECTS OF A CLASS OF SUBJECTIVE GAMES OF INCOMPLETE INFORMATION	E-0

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



EXECUTIVE SUMMARY

The outcome of armed conflict always depends on the decisions made by the participants as the conflict unfolds. As part of its program in mathematics, the Office of Naval Research (ONR) has sponsored work to develop novel models of decisionmaking during conflict, in order to broaden understanding of the factors leading to alternative outcomes. In this final report, prepared under contract N00014-84-C-0458 for ONR, new results in decision modeling and new results obtained using these models are described.

The Navy needs a predictive mathematical theory of decisionmaking in dynamic, multiperson decision situations to better understand command and control in Naval conflict. Models representing both friendly and hostile forces, both sides distributed over the sea with limited communications among participants on each side, must ultimately be treated by formal mathematical theory to provide a sound basis for future system design. The theory should combine results from the mathematics of multiperson decisionmaking with results from behavioral theories of human decisionmaking under uncertainty. This novel combination of approaches would provide valuable qualitative and quantitative insights into the behavior of distributed decisionmakers in complex situations of interest to the Navy. The work reported here is a contribution towards the multidisciplinary theoretical extension needed, drawing on existing prescriptive mathematical approaches to multiperson decisionmaking and on descriptive approaches to human behavior modeling reported in the behavioral decision theory literature.

Mathematical game theory studies decisionmaking in problems with multiple decisionmakers. Game theoretic results characterize decision strategies and outcomes which are rational in the context of specific quantitative behavioral norms. Hence, the purpose of game theory is to prescribe "rational behavior" in a multiperson decision situation. However, game theory has several shortcomings in providing a predictive theory of multiperson decisionmaking:

1. It is based on the assumption that all decisions will be strictly rational decisions, allowing for no deviations from rationality.
2. Rational behavior is defined in terms of Von Neumann's maximum expected utility paradigm [1].
3. It assumes that every decisionmaker has a common representation of the overall decision problem (the rules of the game and the game parameters are common knowledge [2]).
4. It assumes that players are fully-committed to select future decisions in accordance with apriori-selected strategies. It does not permit adaptive decisionmaking behavior.
5. It assumes that processing of information by the decisionmakers will be done optimally, in a Bayesian framework [3].
6. It does not explicitly consider human limitations and behavioral trends in information processing, option evaluation, and action selection.

Empirical research on human decisionmaking [4],[5], has established that all of the above assumptions are systematically violated in specific situations. This work has led to several extensions of the basic framework of game theory. For example, the recent research in perfect, proper and sequential equilibria in nonzero sum games [6]-[8] is concerned with defining rational behavior in a way which takes into account the possibility of "irrational" actions. The work of Aumann and Maschler [9], Ho [10], and others [11] studies games where the assumption of prior commitment to apriori strategies was relaxed. The

work of Harsanyi [12],[13], and others ([15]-[16]) or games of incomplete is aimed at relaxing the assumption that every decisionmaker has a common representation of the overall problem. The work of Kadane and Larkey [17],[18] and Wilson [19] advocates viewing the multiperson decision problem as a set of subjective single person decision problems, where each participant models the expected actions of the other decisionmakers by subjective probabilities.

The purpose of our research effort was to further develop the mathematical theories of multiperson decisionmaking by explicitly considering multiperson decision models which incorporate human limitations and behavioral trends in information processing, option evaluation, and action selection. It is our opinion that a predictive, quantitative theory of multiperson decisionmaking must adopt the viewpoint of a person evaluating his decision alternatives in a real-time, dynamically-evolving situation. Thus, at any time, a human's decision problem separates into three stages:

1. How he interprets the information he's observed in the past,
2. How he ranks his possible decisions at the present time, and
3. How he selects a decision based on these rankings.

These stages correspond to the last three stages in the SHOR paradigm for human decisionmaking discussed in Wohl [20]. The SHOR paradigm is illustrated Fig. 1. Using this paradigm, we were able to organize the classes of human behavioral traits which should be considered in the multiperson decision models. In the first stage, we had to consider biases and limitations in the interpretation and combination of information obtained through observations of the decision problem. In the second stage, we had to consider how individual decisionmakers would evaluate their choices; behavioral decision theory has proposed a number of competing axiomatic models (e.g., [21]) for this process.

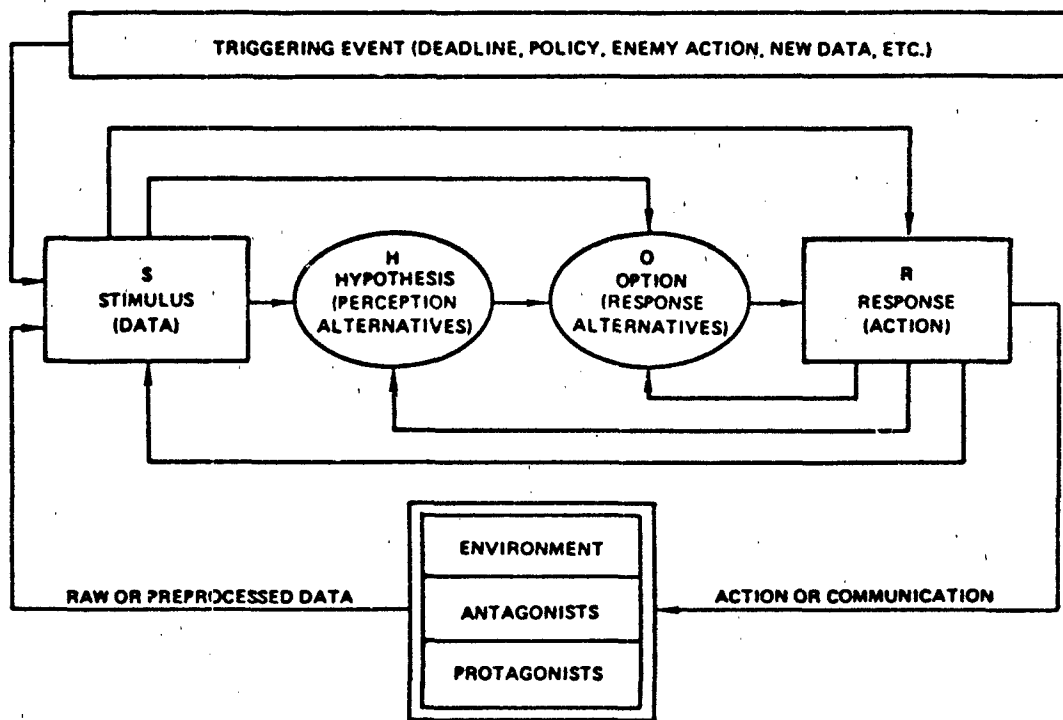


Figure 1. Dynamics of Tactical Decision Process - the SHOR Model

In the last stage, we had to consider how actions were selected based on their evaluation (e.g., [22],[23]).

Starting from this premise, we began our research by reviewing available results on human behavior in information processing, option evaluation and response selection. A brief summary of some of these results is included in Appendix A. Essentially, the literature contains ample evidence that one of the fundamental assumptions of game theory, defining rational behavior in terms of maximization of expected utility, is systematically violated (see Machina's survey [24]). Thus, many alternative formulations for prescribing rational behavior or describing human behavior have been proposed ([4],[5]). Researchers such as Kadane and Larkey [17] have gone as far as suggesting that a prescriptive theory for human decisionmaking in multiperson decision problems should adopt a single decisionmaker perspective, where the actions of other decisionmakers are modeled by a subjective probability over the possible set of actions. There has been considerable debate on this approach, centering on whether the subjective probability should depend on the current choice of action of the decisionmaker, and on how such subjective probabilities are computed [25]-[28]).

The next step in our research addressed the key philosophical question concerning any theory of multiperson decisionmaking, namely: What information does each decisionmaker have concerning the behavioral characteristics of the other decisionmakers? To illustrate the importance of this question, consider the game-theoretic model. In game theory, this information is part of the rules of the game. Hence, this information is represented as common knowledge. Specifically, the utility function of each decisionmaker is known to every other decisionmaker. Using Harsanyi's theory of games of incomplete

information [12], this assumption can be relaxed so that an individual's utility function is known only probabilistically to other decisionmakers. However, this probability distribution is again common knowledge; this means, for example, that decisionmaker 1 has perfect knowledge of how decisionmaker 2 models decisionmaker 1's behavior. Thus, in Harsanyi's theory of games of incomplete information, the assumption of common knowledge has been moved one level higher, to common knowledge of a probability distribution rather than of a specific value. However, this assumption still implies that each decisionmaker knows very well the thinking process of other decisionmakers.

Given the multiple possibilities for models of human decisionmaking described in Appendix A, we felt that the specific parameters of an individual decisionmaker would be known primarily to himself, and not to the other decisionmakers. There are four possible classes of approaches for mathematically modeling the information which each decisionmaker has concerning the behavioral characteristics of the other decisionmakers. These are:

1. Common knowledge,
2. Private knowledge with imperfect information,
3. Private knowledge with incomplete information, and
4. Secret knowledge.

Modeling behavioral characteristics information as common knowledge assumes that each decisionmaker knows every other decisionmaker's behavioral characteristics, and this information is common knowledge in the sense of [2].

Modeling the information as private knowledge with imperfect information means that there is a joint probability distribution over all the decisionmakers' possible behavioral characteristics. This probability distribution is itself common knowledge among decisionmakers; in addition, each decisionmaker is

provided with partial information concerning the behavioral characteristics of the other decisionmakers.

Modeling the information as private knowledge with incomplete information is similar, except that there is no overall joint probability distribution which is common knowledge; in this case, each decisionmaker must subjectively construct this probability distribution. Modeling this information as secret information means that each decisionmaker has a subjective model of every other decisionmaker, and any differences among the models held by different decisionmakers is secret knowledge.

From a mathematical perspective, modeling human characteristics of decisionmakers as common knowledge is the approach which is most akin to game theory. Essentially, human characteristics of decisionmakers would be incorporated into the rules of the game. For example, risk-averse behavior [24] would be represented as a factor in a decisionmaker's utility; similarly, any biases in information processing would be incorporated in the rules for updating of probabilities. Such modification to the expected utility paradigm would result in the violation of several important game theory results. For example, Von Neumann's normalization principle [1], which states that a game in extensive form can be reduced to an equivalent game in normal form, may not be applicable if probabilities do not evolve according to Bayes' rule or if strategies are not selected according to the expected utility model.

Although convenient from a mathematical perspective, the common knowledge approach is the least satisfactory from a modeling perspective, because it assumes that each decisionmaker has a very accurate model of the other decisionmakers. The private knowledge with imperfect information approach makes

weaker assumptions concerning the knowledge provided to decisionmakers concerning the knowledge provided to decisionmakers concerning the behavior of other decisionmakers. However, the imperfect knowledge is modeled by a probability distribution which is common knowledge. It is not clear that in multiperson decision situations, such a state of common knowledge can be achieved without an extensive cooperative bargaining session to agree on this probability. For many military situations of interest, this may not be possible.

From a practical perspective, the last two approaches capture the subjective nature of human decisionmaking best. In the private knowledge with incomplete information approach, each decisionmaker can subjectively estimate a probability distribution for the decisionmaking characteristics of every other decisionmaker. However, these estimates are subjective estimates, and there is no reason that they must be consistent across decisionmakers. In the secret knowledge approach, each decisionmaker models the human characteristics of other decisionmakers as values rather than probability distributions and assumes that these values are correct. Again, the fact that these values may be incorrect is secret knowledge to each decisionmaker. In both of these approaches, each decisionmaker has a subjective model of how he and other decisionmakers will make decisions. However, these models need not be consistent across decisionmakers. This leads to a number of interesting questions which we addressed in our research.

1. Do qualitative results characterizing outcomes of multiperson decision situations change due to the subjective model differences among decisionmakers?
2. How do decisionmakers interpret information they receive from other decisionmakers?
3. How do they incorporate this information to form their decisions?

4. Do the decisionmakers realize during the play of the game that their decision models are inconsistent?
5. Do those inconsistencies prevent the decisionmakers from reaching a desirable outcome?
6. When inconsistencies are detected, how do the decisionmakers modify their subjective models?

In order to study these questions in a specific context, we considered two classes of multiperson decisionmaking problems:

1. Consensus Problems, and
2. Two-Person nonzero sum games of incomplete information.

In consensus problems, multiple decisionmakers with private information and a common goal are trying to reach a consensus decision by proposing tentative decisions recursively among themselves. Consensus problems are a simple class of cooperative decision problems emphasizing the implicit information transfer among decisionmakers through the choice of decisions (signaling). Under simple communication conditions, the results of [29]-[34] establish that, in the absence of subjective model differences, a consensus is always reached. However, differences in subjective models among decisionmakers may lead to misunderstandings in the signaling processes, thereby preventing the decisionmakers from reaching a consensus.

Our first investigation of these problems is summarized in the paper in Appendix B. In this paper, we show that, for a specific consensus problem of estimating the probability of an event, when the individual differences among decisionmakers are secret knowledge, the consensus process can reach a state of contradiction. In this state, the signals from one decisionmaker cannot be interpreted by another decisionmaker due to the secret difference between their models. Thus, the secret information that the subjective models were

different becomes common knowledge. In Appendix C, we show that this phenomenon is typical of general consensus problems when differences in subjective decisionmakers are secret knowledge. Furthermore, we establish that there are only two classes of qualitative outcomes possible: a consensus is reached, so the differences in subjective models remain secret knowledge, or a contradiction is reached, so the differences in subjective models become common knowledge.

One question which was not addressed in the results in Appendices B and C concerned the likelihood of the two classes of outcomes. For specific problems, how likely was a contradictory outcome? In Appendix D, we develop a number of results which provide answers to this question. Specifically, we derive conditions, depending on the type of consensus problem, which characterize whether the set of subjective models which result in contradictions is a dense set in the set of possible models, using a specific topology. We also derive conditions which characterize when the consensus process is robust to small differences in subjective models.

In the paper in Appendix D, we also study the consensus problem when the differences among subjective models are represented as private information with either incomplete or imperfect knowledge. That is, there is a second-order probability model which describes each decisionmaker's subjective beliefs concerning the decision model of the other decisionmakers. The results in Appendix D establish that the two classes of qualitative outcomes described previously can occur. In addition, there is yet a third class of outcomes, which can be described as follows. In these outcomes, the consensus process reaches a state where each decisionmaker cannot learn any additional

information from the tentative decisions communicated by other decisionmakers. However, there is a significant difference in information among decisionmakers, so that a consensus is not reached.

The result in Appendix D is very surprising. Consider a consensus process where two decisionmakers have exactly the same decision model. According to the results of [32], a consensus will be reached eventually. However, assume now that both decisionmakers have exactly the same decision model, but they do not know it for certain! The results in Appendix D show that, in this situation, this lack of certainty can lead the consensus process to a stalled state, from which a consensus cannot be reached. In essence, the lack of certainty prevents the consensus process from transmitting enough information to achieve a consensus outcome.

This result has important implications for more general classes of multi-person decision problems, such as cooperative bargaining problems. Loosely interpreted, it states that the presence of uncertainty concerning behavioral characteristics of the bargaining decisionmakers can lead to a stalemate in the bargaining process. In practical situations, this stalemate can be broken by changing the nature of the bargaining process, such as bringing in an arbitrator.

In the area of nonzero sum games of incomplete information, we considered a simple class of two-person symmetric games. We developed an axiomatic formulation describing the subjective decision model approach discussed above. Based on this formulation, we studied the properties of the rational strategies and outcomes of both the static and infinitely-repeated versions of the game. The results in Appendix E show that, under the assumption that each

decisionmaker is subjectively rational according to his own internal model* in the static case, the existence of differences between subjective models becomes common knowledge. However, in the infinitely-repeated game, these differences are ameliorated during the play of the game, so that there exist core equilibrium strategies which are rational from each decisionmaker's perspective. In Appendix E, we show that whether the subjective differences become common knowledge or cannot depend on the specific bargaining strategy used for selecting among potential core equilibrium outcomes.

In sum, our research has investigated two classes of mathematical models of multiperson decisionmaking in dynamic situations which focus on the subjective nature of human behavioral characteristics in information processing, option evaluation and action selection. Our mathematical formulation is similar in spirit to Krdane and Larkey's [17], in that individual rationality is defined in terms of subjective models. However, we do not propose to replace the game-theoretic models in these subjective models by decision theory models using subjective probabilities of action for other decisionmakers. Instead, we have developed a mathematical formulation which allows for differences in subjective game theory models among decisionmakers. Our research focused on studying the qualitative implications of these differences in cooperative and noncooperative paradigms. The results developed in the appendices to this report shed considerable light on these implications in dynamic decisionmaking under uncertainty.

Nevertheless, many technical problems remain to be addressed. Perhaps the key question remaining unresolved is how to effect the resolution of model

*For an interesting discussion of the application of internal models to representing Navy Command and Control decisionmaking, see Athans [35].

differences once they are discovered. The approach proposed in Appendix D requires that an inference problem be started at that time to properly identify the model differences. This approach does not indicate how to resolve these differences to achieve a desirable qualitative outcome, either through a common calibration process or a bargaining process.

The subjective game framework which was used throughout our research provides a foundation for the development of a quantitative predictive theory of multiperson decisionmaking. In our opinion, the next major step is the development of specific behavioral theories and mathematical models of these theories for human decisionmaking in dynamic multiperson decision situations. As seen in the survey in Appendix A, the state of the research in individual decision theory has focused on developing alternatives to the expected utility paradigm [1]; the numerous competing alternatives must be narrowed down and shown to be superior to the expected utility paradigm. Furthermore, additional research is needed to focus on the multiperson aspects of behavioral theories of information processing and option selection. Based on these results, specific quantitative models of multiperson decisionmaking can be developed and evaluated in empirical research. These research directions are being followed in many research programs, so that progress towards the development of a quantitative predictive theory of multiperson decisionmaking will continue.

APPENDIX A

AN OVERVIEW OF RESULTS OF HUMAN DECISIONMAKING

A-0

APPENDIX A

AN OVERVIEW OF RESULTS IN HUMAN DECISIONMAKING

A.1 INTRODUCTION

In this appendix, we overview some of the principal results available in the literature concerning empirical studies of how humans make decisions. We have organized these results into two sections: Human Inference, dealing with the problem of situation assessment and hypothesis evaluations, and utility theory, dealing with the problems of option evaluation and option selection. For a more detailed review of these disciplines the reader is referred to the excellent surveys and books by Rapoport and Wallsten [1972], Slovic, Fishhoff and Lichtenstein [1977], Einhorn and Hogarth [1981], Kahneman, Slovic and Tversky [1982], Schoemaker [1982], Fishburn [1982] and Machina [1983].

A.2 HUMAN BEHAVIOR IN SITUATION ASSESSMENT

Since most of the decision tasks involve uncertainty, considerable effort has been spent in studying how people formulate and change their opinions about uncertain hypotheses. The literature in the area of subjective probability assessment and revision of opinion shows two different approaches to the modeling problem. The first approach, advanced by statisticians and psychologists, is based on probability theory and statistics and relies on the concept of a "statistical man" - an optimal, Bayesian inferer (observer). Bayes' rule provides a normative representation of how a decisionmaker should revise his probability estimates on the basis of new information. The

descriptive considerations have been handled primarily by adjusting the functional form of the normative model. This approach led to the study of "conservatism" - a suboptimal human behavior that produces posterior probabilities nearer to the prior probabilities than those specified by Bayes' rule (Edwards and Phillips [1964]; Peterson and Beach [1967]). The second approach, proposed mainly by psychologists, argues that the human is a selective, sequential information processor with limited capacity (Hogarth [1975]). It is hypothesized that this limited information-processing capacity leads him to apply simple heuristics and cognitive strategies which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. Much of the work on this judgmental heuristics has been performed by Tversky and Kahneman [1971], [1973], [1974]. They demonstrated that three judgmental heuristics - representativeness, availability, adjustments and anchoring - determine probabilistic inferences in many tasks. However, these findings can only be described in qualitative terms and, as yet, no quantitative descriptive theory based on the heuristics has emerged.

A.2.1 Bayesian Revision of Opinion

Bayes' rule has provided much of the impetus to the research on normative-descriptive modeling of judgmental processes. A basic hypothesis is that opinions (judgments) should be expressed in terms of subjective probabilities and that the optimal revision of such opinions must be accomplished via Bayes' rule. A considerable number of studies, involving mainly binomial (so called "bookbag and poker chip" games) and multinomial tasks, have compared subjects' numerical probability assessments to those predicted by Bayes' rule (Edwards and Phillips [1964]; Peterson and Beach [1967], Donnel and Duchrame [1975]). A general, but by no means universal, conclusion has been

that the estimates were monotonically related to those specified by Bayes' rule, but were conservative. That is, the posterior probabilities estimated by subjects were nearer to the prior probabilities than those obtained via Bayes' rule. Several explanations are offered for the phenomena of conservatism. It is believed that conservatism is due, in small part, to procedural variables (e.g., payoffs and incentives, sample size, sequential ordering of the data, prior probabilities, etc.) and, in large part, to subjects' misperception of the underlying sampling distributions, misaggregation of the data, or simply response bias. Misperception is generally attributed to the mismatch between subjective and actual (objective) probability distributions (Peterson, Duchrane and Edwards [1968]; Wheeler and Beach [1968], Lichtenstein and Feeney [1968]), and to the human tendency to discount the importance of rare events when they occur (Vlek and Van der Heijden [1967]). Misaggregation refers to the nonoptimal sequential revision of subjective probabilities and has been advanced as the major source of conservatism primarily by Edwards and his associates (Edwards [1968], Edwards, Phillips, Hayes and Goodman [1968]). The notion of response bias was advanced by Peterson [1968] and is related to subjects' unwillingness to use extreme numbers and odds. A comprehensive review of the issue of conservatism is provided by Slovic and Lichtenstein [1971], Rapoport and Wallsten [1972] and Slovic, Fishoff and Lichtenstein [1977].

Along different lines, Kahneman and Tversky [1973], Tversky and Kahneman [1974] and Grether [1980] among others, have found that probability updating underweighs prior information and overweighs the representativeness of the current sample. This phenomena is similar to Tversky and Kahneman's [1971]

law of small numbers. Tversky and Kahneman [1983] have also found biases in the combination of evidence from independent and correlated sources.

In an attempt to overcome the descriptive deficiencies of the Bayesian model, several empirical modifications to Bayes' rule have been offered. These modifications can be embedded into a generalized version of Bayes' rule, where an additional term is multiplied to the likelihood of each hypothesis. This term is called a disability or impediment function in Edwards and Phillips [1964], and is supposed to capture the suboptimal nature of human information processing (Snapper and Frybach [1971]).

Although Bayesian revision of opinion can be studied as a separate phenomena, it is most useful when interwoven with decisionmaking and action selection. The posterior probabilities of various hypotheses (states of nature) can be used, in combination with information about payoffs associated with various decisions and states of nature, to maximize the (subjective) expected value, the (subjective) expected utility, or whatever criterion of optimality. Human performance modeling in signal-detection tasks exemplifies this approach. The task of the subject is to decide whether or not a signal is present in a block of observations.

Signal-detection experiments have been conducted in a wide variety of contexts. Examples are experiments in sound localization (Voelcker [1961]), detection of movement (Kinchela and Allen [1969]), speech recognition (Egan, et al. [1961]), and recognition of memorized words (Parks [1966]). A comprehensive exposition of signal-detection theory and psychophysics has been provided by Green and Swets [1966], and a fine summary of the theory is provided by Sheridan and Farrel [1974].

The experimental results show that the human performance is monotonically related to those predicted by the model. It is possible to manipulate subjective decision thresholds (criteria) by varying prior probabilities and pay-offs. However, the amount of change has been found to be less than optimal. The subjects also have difficulty aggregating information across a sequence of trials (Swets and Green [1961]) - a tendency similar to conservatism in Bayesian revision of opinion.

An additional effect in human inference which cannot be represented using impediment functions is the asymmetry between effect-cause inferences and cause-effect inferences noted in Ajzen [1977], Tversky and Kahneman [1980a] and Einhorn and Hogarth [1981]. The empirical evidence indicates that information which receives a causal interpretation is weighed more heavily in judgment than information that is diagnostic. Tversky and Kahneman [1980] correlate their results with previous research (Janis [1972]), (Jervis [1975]) indicating that humans overestimate the accuracy of uncertain models in predicting behavior, and, when confronted with evidence concerning the errors in their models, would rather find a plausible explanation than revise their models. These results are particularly relevant in multiperson decision-making, where the interactions among decisionmakers force each decisionmaker to develop internal models of the other decisionmakers.

A.2.2 Judgmental Heuristics and Biases

Recent research on probabilistic judgments has focused on the discovery and description of heuristics, or simple cognitive strategies, that are employed in the quantification of uncertainty. Much of this work has been performed by Tversky and Kahneman [1971], [1973], [1974]. Their research centers

on the determination of how people evaluate uncertainty rather than how well they evaluate it. Tversky and Kahneman demonstrated, qualitatively, that three judgmental heuristics - representativeness, availability, adjustment and anchoring - determine probabilistic judgments in a variety of tasks. These qualitative findings are potentially useful in the development of a (as yet elusive) quantitative, descriptive theory of judgment.

1. REPRESENTATIVENESS HYPOTHESES

Tversky and Kahneman [1972] hypothesized that people evaluate the probability of an event on the basis of the degree of similarity between the event and the evidence they have examined. If the degree of similarity is high, then the probability of the event is judged to be high. It was demonstrated that the representativeness heuristic can explain people's intuitive predictions that were at variance with the normative judgments. This was accomplished by showing the insensitivity of the representativeness heuristic to several normatively important factors of judgment, viz., the representativeness heuristic is liable to be used when the general properties of events are emphasized.

11. AVAILABILITY HEURISTIC

With this heuristic, people evaluate the probability of an event on the basis of the ease with which instances or occurrences can be recalled or imagined. Availability is a valid cue for the assessment of probability because, in general, instances of more frequent events are recalled more easily than the instances of less frequent events. However, availability is also affected by other factors unrelated to probability. Consequently, availability heuristic results have been used in systematic biases, some of which follow:

- (1) Biases due to retrievability of instances: An event whose instances are easily recalled will appear more frequently than an event of equal probability (or frequently) whose instances are less easily recalled.
- (2) Biases due to the effectiveness of a search set: In tasks requiring the estimates of the relative frequencies of words, the availability heuristic leads to a judgment that the frequency of occurrence of abstract words (e.g., love in love stories) is much higher than the concrete words (e.g., door).
- (3) Biases of imaginability: In tasks in which one must assess the frequency of an event whose instances are not stored in memory, one may generate the instances according to an algorithm or a rule. The ease with which one can generate instances forms the basis for probability or frequency assessment. Depending on the nature of the rule, this mode of probability assessment may lead to serious biases.
- (4) Illusory correlation: The probability of how frequently two events co-occur is related to the associative bond between them. Strong associates will be judged to have occurred frequently together.

It is stated that the availability heuristic is used when events are thought of in terms of specific instances.

111. ADJUSTMENT AND ANCHORING

With this heuristic, an initial value or anchor is used as a first approximation to the judgment. The initial value is then adjusted according to the information provided. Typically, these adjustments are imperfect and insufficient. Even payoffs for accuracy did not reduce this effect. It was demonstrated that anchoring could explain people's judgment of the probabilities of conjunctive and disjunctive events and their assessment of the variability of probability distributions.

Studies of decisions under risk indicate that people tend to overestimate the probability of conjunctive events and underestimate the probability of

disjunctive events. With regard to the variability of probability distributions, several investigators found that subjects state more narrow confidence intervals than are justified by the evidence presented to them. Edwards [1975] noted that this peculiarity is dependent entirely on the format of questions.

A comprehensive review of heuristics in probabilistic judgment has been provided by Slovic, Fishoff and Lichtenstein [1977]. Although the evidence suggests that heuristics are employed in the assessment of uncertainty, the specific heuristic selected, the way it is used, and the quality of judgment it provides, are all highly problem dependent. Therefore, heuristics may be thought of as explanatory psychological processes of human probabilistic judgment and cannot be regarded as a general theory of judgment.

A.3 EMPIRICAL RESEARCH IN UTILITY THEORY

As stated in Machina's survey [1983], the theoretical underpinnings of single person decision theory under uncertainty are based on the expected utility hypothesis of individual behavior. This hypothesis essentially states that, when faced with alternative risky prospects over a set of outcomes, a rational decisionmaker will always choose a prospect which yields the highest expectation of some utility function defined over the set of outcomes. This utility function, often called a Von Neumann-Morgenstern utility function, represents the preference of individuals over the different outcomes.

As noted in Fishburn [1982], the existence of Von Neumann-Morgenstern utility function and the validity of the expected utility hypothesis follow from some variations of the following three axioms:

- (1) Transitivity of preferences over prospects; i.e., if prospect A is weakly preferred to prospect B, and prospect B is weakly preferred to prospect C, then prospect A is weakly preferred to prospect C.
- (2) Continuity of preferences over prospects; i.e., if two sequences of prospects A_n , B_n are such that A_n is weakly preferred to B_n , and A_n , B_n converge to A, B respectively (in the topology of weak convergence), then A is weakly preferred to B.
- (3) Independence of preferences over common alternatives; i.e., if A is weakly preferred to B, then, for any C, and any $0 < \alpha < 1$, $\alpha A + (1-\alpha)C$ is weakly preferred to $\alpha B + (1-\alpha)C$.

These axioms appear to be simple conditions which the preferences of any rational decisionmaker should satisfy. Hence, empirical results concerning the validity of these axioms have been limited. As Machina states: "Unfortunately, though no doubt due in part to the widely held belief in the inherent 'rationality' of the expected utility axioms and in part to the tremendous success of the theoretical developments during this period, the last few decades have seen nowhere near the amount of empirical estimation and testing by economists that such a widely used model of behavior ought to have received." In this section, we will overview some of the available results on empirical research in utility theory.

One of the first empirical studies concerned the risk behavior of individuals, and their desire both to participate in high-stakes risk-seeking lotteries, and risk-averse insurance purchases. Friedman and Savage [1948] postulated that the utility function should be a function that is concave, "locally risk averse", about low wealth levels, and convex (risk seeking) at high wealth levels. Hence, an individual whose current wealth position was near the inflection point would indeed purchase both insurance against losses and lottery tickets offering a small chance of large gains. This utility function implies two other commonly observed aspects of individual preferences

over uncertain prospects. First, individuals prefer increases in risk in the upper tails of already random wealth distributions over risk increases in the lower tails of such distributions. Second, individuals prefer distributions with large right tails over large left tails (Kraus and Litzenberger [1976], Scott and Horvath [1980]).

There is evidence to suggest that the Friedman-Savage characterization of the utility functions (risk-averse over biases, risk-seeking over gains) is incomplete. Among others, Hershey and Schoemaker [1980a], [1980b] and Kahneman and Tversky [1979] have found both risk-averse and risk-seeking behavior in loss situations. Kahneman and Tversky [1979] attribute this observed behavior to nonlinearities in the probabilities. Hershey and Schoemaker [1980a] also found that problem representations affected the decisionmaking behavior. The differences in behavior introduced by problem representation were observed in many other studies (Slovic [1969a], Schoemaker and Kunreuther [1979], Tversky and Kahneman [1981], Kahneman and Tversky [1982]) and led to the theory of "framing," where the mental point for defining what is a gain versus what is a loss depends on the specific wording of the decision problem.

There have been a number of empirical studies to test the validity of the axioms of expected utility theory. The most critical of these axioms is the independence axiom. As noted in Fishburn [1982] the independence axiom is the primary normative assertion of expected utility theory, because it restricts the expected utility to be a linear functional of the probability distribution over the set of outcomes. However, examples such as Allais' paradox [1953] have produced evidence that this axiom is violated often (Allais and Hagen [1979], Raiffa [1968], Slovic and Tversky [1974], Kahneman and Tversky [1979], MacCrimmon and Larsson [1979]).

The assumption of transitive preferences has also been tested empirically. Early studies in utility theory (Arrow [1951]) indicated that in many choice situations, preferences were not transitive. Empirical research by Edwards [1954], Weinstein [1968] and Tversky [1969] have produced examples of intransitive preferences among alternatives. In subsequent research, Lichtenstein and Slovic [1973] have found that, when faced with a direct choice of prospects A and B, decisionmakers prefer A to B even though they assign a higher certainty equivalent value to B than to A.

Another systematic violation of the expected utility hypothesis arises in the context of choice. Specifically, the expected utility hypothesis predicts that the prospect which maximizes expected utility will always be selected. However, when repeatedly confronted with the same pair of prospects, worded in the same manner, individuals will not always make the same choice. The early results of Mosteller and Nogee [1951] established that the individual's choice probabilities were continuous, monotonic functions of the differences in the expected utilities of the prospects. Explicit models of randomized choice have been proposed by Luce [1959] and Fishburn [1976], [1978] among others.

At a more fundamental level, the expected utility hypothesis postulates the existence of probabilities which obey the standard calculus (Savage [1954]). However, Ellsberg's paradox (Ellsberg [1961]) demonstrates that rational decisionmakers may weigh the alternatives with scale factors which do not obey the laws of probability. Ellsberg's paradox has been verified experimentally by MacCrimmon [1965], Slovic and Tversky [1974], and MacCrimmon and Larsson [1979]. Plausible explanations for this behavior have been offered by Arrow and Hurwicz [1972] in terms of maxmin choices instead of maximum expected utility.

A.4 CONCLUSION

This appendix has provided a brief overview of some results available concerning human decisionmaking. The focus has been to study results which indicate that the standard normative theories such as probability theory and utility theory fall short of describing observed decisionmaking behavior. The rest of the research in this report has been motivated by the existence of these differences.

REFERENCES

- Allais, M. [1953], "Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Americaine," Econometrica, 21, pp. 503-546.
- Allais, M. and O. Hagen [1979] (eds.), Expected Utility Hypotheses and the Allais Paradox, Dordrecht, Holland: D. Reidel Publishing Co.
- Ajzen, I. [1977], "Intuitive Theories of Events and the Effects of Base-rate Information on Prediction," J. Pers. Soc. Psychol., 35, pp. 303-314.
- Arrow, K. [1951a], "Alternative Approaches to the Theory of Choice in Risk-Taking Situations," Econometrica, 19, pp. 404-437.
- Arrow, K. [1951b], Social Choice and Individual Values, New Haven: Yale University Press.
- Arrow, K. and L. Hurwicz [1972], "An Optimality Criterion for Decision-Making under Ignorance," in ().
- Donnel, M.L., and W.M. Duchrame [1975], "The Effect of Bayesian Feedback on Learning in an Odds Estimation Tasks," Organizational Behavior and Human Performance, 14, pp. 305-313.
- Edwards, W. [1954], "Probability-Preferences Among Bets With Differing Expected Value," American Journal of Psychology, 67, pp. 56-67.
- Edwards, W. [1968], "Conservatism in Human Information Processing," in B. Kleinmuntz (ed.), Formal Representations of Human Judgement, pp. 17-52, New York: Wiley.
- Edwards, W. [1975], "Comments on Cognitive Processes and Probability Assessment, by R.M. Hogarth," Journal of American Statistical Society, 70, pp. 291-293.
- Edwards, W. and L.D. Phillips [1964], "Man as Transducer for Probabilities in Bayesian Command and Control Systems," in M.W. Shelly, II and G.L. Bryan (eds.) Human Judgements and Optimality, pp. 360-401, New York: Wiley.
- Edwards, W., L.D. Phillips, W.L. Hayes and B.C. Goodman [1968], "Probabilistic Information Processing Systems: Design and Evaluation," IEEE Trans. on Systems Science and Cybernetics, 4, pp. 248-265.

- Egan, J.P., F.R. Clark, and E.C. Carterette [1961], "On The Transmission and Confirmation of Messages in Noise," Journal of Acoustical Society of America, 33, pp. 771-778.
- Einhorn, H. and R. Hogarth [1981], "Behavioral Decision Theory: Processes of Judgment and Choice," Annual Review of Psychology, 32, pp. 53-88.
- Ellsberg, D. [1961], "Risk, Ambiguity and the Savage Axioms," Quarterly Journal of Economics, 75, pp. 643-669.
- Fishburn, P. [1976], "Unbounded Utility Functions in Expected Utility Theory," Quarterly Journal of Economics, 90, pp. 163-168.
- Fishburn, P. [1978], "On Handa's 'New Theory of Cardinal Utility' and the Maximization of Expected Return," Journal of Political Economy, 86, pp. 321-324.
- Fishburn, P. [1982], The Foundations of Expected Utility, Dordrecht, Holland: D. Reidel Publishing Co.
- Friedman, M. and L. Savage [1948], "The Utility Analysis of Choices Involving Risk," Journal of Political Economy, 56, pp. 279-304.
- Green, D.M. and J.A. Swets [1966], Signal Detection Theory and Psychophysics, New York: Wiley.
- Grether, D. [1980], "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," Quarterly Journal of Economics, 95, pp. 537-557.
- Hershey and Schoemaker [1980a], "Risk-Taking and Problem Context in the Domain of Losses — An Expected Utility Analysis," Journal of Risk and Insurance, 47, pp. 111-132.
- Hershey and Schoemaker [1980b], "Prospect Theory's Reflection Hypothesis: A Critical Examination," Organizational Behavior and Human Performance, 25, pp. 395-418.
- Hogarth, R.M. [1975], "Cognitive Processes and the Assessment of Subjective Probability Distributions," Journal of the American Statistical Association, 70, pp. 271-289.
- Janis, I.L. [1972], Victims of Groupthink, Boston: Houghton-Mifflin.
- Jervis, R. [1975], Perception and Misperception in International Relations, Princeton: Princeton University Press.
- Kahneman, D. and A. Tversky [1972], "Subjective Probability: A Judgement of Representativeness," Cognitive Psychology, 3, 430-454.

- Kahneman, D. and A. Tversky [1973], "On the Psychology of Prediction," Psychological Review, 80, pp. 237-251. Reprinted in Kahneman, Slovic and Tversky [1982].
- Kahneman, D. and A. Tversky [1979], "Prospect Theory: An Analysis of Decision Under Risk," Econometrica, 47, pp. 263-291.
- Kahneman, D. and A. Tversky [1982], "The Psychology of Preferences," Scientific American, 246, pp. 160-173.
- Kahneman, D. P. Slovic and A. Tversky [1982], Judgement Under Uncertainty: Heuristics and Biases, Cambridge: Cambridge University Press.
- Kinchela, R.A. and L.G. Allen [1969], "A Theory of Visual Movement Perception," Psychological Review, 76, pp. 537-558.
- Kraus, A. and R. Litzenberger [1976], "Skewness Preference and the Valuation Risk Assets," Journal of Finance, 31, pp. 1085-1099.
- Lichtenstein, S. and G.J. Feeney [1968], "The Importance of The Data Generating Model on Probability Estimating," Organizational Behavior and Human Performance, 3, pp. 62-67.
- Lichtenstein, S. and P. Slovic [1973], "Response-Induced Reversals of Preference in Gambling: An Extended Replication in Las Vegas," Journal of Experimental Psychology, 101, pp. 16-20.
- Luce, R. [1959], Individual Choice Behavior: A Theoretical Analysis, New York: John Wiley and Sons.
- MacCrimmon, K. [1965], An Experimental Study of the Decision Making Behavior of Business Executives, Doctoral Dissertation, University of California, Los Angeles.
- MacCrimmon, K and S. Larsson [1979], "Utility Theory: Axioms Versus 'Paradoxes'," in Allais and Hagen [1979].
- Machina, M. [1983a], "Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom," in Stigum and Wenstop [1983].
- Machina, M. [1983b], "The Economic Theory of Individual Behavior Toward Risk: Theory, Evidence and New Directions," IMSSS Report 433, Stanford University, October 1983.
- Mosteller, F. and P. Nogee [1951], "An Experimental Measurement of Utility," Journal of Political Economy, 59, pp. 371-404.
- Parks, T.E. [1966], "Signal-Detectability Theory of Recognition Memory Performance," Psychological Review, 73, pp. 44-58.

- Peterson, C.R. [1968], "Aggregating Information About Signal and Noise," Proceedings 76th Annual Convention APA, 3, pp. 123-124.
- Peterson, C.R. and L.R. Beach [1967], "Man As An Intuitive Statistician," Psychological Bulletin, 68, pp. 29-46.
- Peterson, C.R., W.M. Duchrame and W. Edwards [1968], "Sampling Distributions and Probability Revisions," Journal of Experimental Psychology, 76, pp. 236-243.
- Raiffa, H. [1968], Decision Analysis: Introductory Lectures on Choice Under Uncertainty, Reading, Massachusetts: Addison-Wesley.
- Rapoport, A. and T. Wallsten [1972], "Individual Decision Behavior," Annual Review of Psychology, 23, pp. 131-176.
- Savage, L. [1954], The Foundations of Statistics, New York: John Wiley and Sons.
- Schoemaker, P. [1982], "The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations," Journal of Economic Literature, 20, pp. 529-563.
- Schoemaker, P. and H. Kunreuther [1979], "An Experimental Study of Insurance Decisions," Journal of Risk and Insurance, 46, pp. 603-618.
- Scott, R. and P. Horvath [1980], "On the Direction of Preference for Moments of Higher Order than the Variance," Journal of Finance, 35, pp. 915-919.
- Sheridan, T.B. and W.R. Farrel [1974], Man-Machine Systems, Cambridge, Massachusetts, M.I.T. Press.
- Slovic, P. [1969], "Manipulating the Attractiveness of a Gamble Without Changing its Expected Value," Journal of Experimental Psychology, 79, pp. 139-145.
- Slovic, P. and S.C. Lichtenstein [1971], "Comparison of Bayesian and Regression Approaches to The Study of Information Processing Judgement," Organizational Behavior and Human Performance, 6, pp. 649-744.
- Slovic, P., B. Fishoff and S. Lichtenstein [1977], "Behavioral Decision Theory," Annual Review of Psychology, 28, pp. 1-39.
- Slovic, P. and A. Tversky [1974], "Who Accepts Savage's Axiom?," Behavioral Science, 19, pp. 368-373.
- Snapper, K.J. and D.G. Frybach [1971], "Inferences Based on Unreliable Reports," Journal of Experimental Psychology, 87, pp. 401-404.
- Stigum, B. and F. Wenstop [1983], Foundations of Utility and Risk Theory with Applications, Dordrecht, Holland: D. Reidel Publishing Co.

- Swets, J.A. and D.M. Green [1961], "Sequential Observations by Human Observers of Signals and Noise," in C. Cherry (ed.), Information Theory, pp. 177-195, London: Butterworth.
- Tversky, A. [1969], "Intransitivity of Preferences," Psychological Review, 76, pp. 31-48.
- Tversky, A. [1972], "Elimination by Aspects: A Theory of Choice," Psychological Review, 79, pp. 281-299.
- Tversky, A. and D. Kahneman [1971], "Belief in the Law of Small Numbers," Psychological Bulletin, 2, pp. 105-110. Reprinted in Kahneman, Slovic and Tversky [1982].
- Tversky, A. and D. Kahneman [1973], "Availability: A Heuristic for Judging Frequency and Probability," Cogn. Psychol., 5, pp. 207-222.
- Tversky, A. and D. Kahneman [1974], "Judgement Under Uncertainty: Heuristics and Biases," Science, 185, pp. 1124-1131.
- Tversky, A. and D. Kahneman [1980], "Casual Schemas in Judgments Under Uncertainty," Progress in Social Psychology, M. Fishbein (ed.), Hillsdale, New Jersey: Erlbaum.
- Tversky, A. and D. Kahneman [1981], "The Framing of Decisions and the Psychology of Choice," Science, 211, pp. 453-458.
- Tversky, A. and D. Kahneman [1983], "Extensional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement," Psychological Review.
- Vlek, C.A.J., and L.H.C. Van der Heijden [1967], "Subjective Likelihood Functions and Variations in the Accuracy of Probabilistic Information Processing," Report E 017-67, Psychol. Institute; Univ. Leiden, The Netherlands.
- Voelcker, H.B. [1961], "A Decision - Theory Approach to Sound Lateralization," in C. Cherry (ed.), Information Theory, London: Butterworth.
- Weinstein, A. [1968], "Individual Preference Intransitivity," Southern Economic Journal, 34, pp. 335-343.
- Wheeler, G. and L.R. Beach [1968], "Subjective Sampling Distributions and Conservatism," Organizational Behavior and Human Performance, 3, pp. 36-46.

APPENDIX B

TECHNICAL REPORT TR-172: ASYMPTOTIC AGREEMENT IN
DISTRIBUTED ESTIMATION WITH INCONSISTENT BELIEFS*

*This research was also supported by DOE contract DEACO-1-80-RA50418

B-0

ASYMPTOTIC AGREEMENT IN DISTRIBUTED ESTIMATION WITH INCONSISTENT BELIEFS

1. INTRODUCTION

Consider two agents, 1 and 2, who wish to estimate the same random variable x . Initially agent i ($i=1,2$) observes the variable y^i ($i=1,2$). Based on his observation Agent $i(2)$ generates an estimate \hat{x}^i (x^2) which he sends to Agent $2(1)$. It is assumed that the message is received without any distortion. Each time an agent receives a message he recomputes a new estimate, based on the original observation* and the messages received by the other agent up until that time, which he then transmits to the other agent. Several questions related to the evolution of these sequences of estimates arise: Will an agent settle on a final estimate? Will the estimates of the two agents eventually agree?

A substantial effort has been recently devoted to the problem of reaching a consensus of opinion among several decisionmakers [1]-[5]**. The crucial assumption in [1]-[5] is the following: All agents are assumed to be Bayesian. Agent i 's view of the world is represented by an a priori distribution p^i on the space of "primitive" random variables x , y^i , and p^i is the same for all agents. Under this assumption the conditions under which asymptotic agreement is achieved are investigated in [1]-[5]. A major result reported in [3] is

*We restrict attention to the case where each agent takes one measurement without any loss of generality; the case where the agents take noise measurements and communicate simultaneously can be analyzed in a similar fashion.

**References are indicated by numbers in square brackets, the list appears at the end of the main body of this report.

the following: if the messages exchanged among the agents are the conditional expectations of the random variable x then the agents agree asymptotically.

In this report we consider two Bayesian agents who have different views of the world and exchange their conditional expectations of the random variable x . We show that asymptotically the two agents either agree or they realize that they have different models and stop communicating any further. Agreement or disagreement depend on the order of communication.

The remainder of this report is organized as follows: The model is presented in Section 2; the process of expectation formation is described in Section 3; the question of convergence of the estimates and of asymptotic agreement are investigated in Section 4. Conclusions are presented in Section 5.

2. THE MODEL

We consider two agents, 1 and 2, and a random variable x which each agent wishes to estimate. We make the following assumptions:

(A1) Both agents are Bayesian. Agent i 's ($i=1,2$) view of the world is represented by an a priori distribution P^i ($i=1,2$) on the space of the "primitive" random variables X, Y^1, Y^2 , whose sample space is Ω . Agent i assumes that the other agent's distribution is also P^i . We do not assume $P^1=P^2$.

(A2) $E^i|x| < \infty$ $i=1,2$ where E^i denotes the expectation with respect to the probability measure p^i induced by P^i .

(A3) Each agent takes only one measurement y^i ($i=1,2$) at time $t=0^*$, and computes the conditional expectation \hat{x}^i of the random variable x based on his measurement. The spaces Y^i ($i=1,2$) are finite.

(A4) At time $t=2K-1$ ($2K$), $K=1,2,\dots$, agent 1(2) sends the message \hat{x}_k^1 (\hat{x}_k^2) to agent 2(1). Thus, messages are transmitted at $t=1,2$, in the following order: $\hat{x}_1^1, \hat{x}_1^2, \hat{x}_2^1, \hat{x}_2^2, \hat{x}_3^1, \hat{x}_3^2, \dots$

The messages \hat{x}_k^i ($i=1,2$, $k=1,2,\dots$) of each agent are the expectations of x based on the original observation Y^i and the messages received in the past by the other agent.

*The problem where the agents take more measurements and communicate simultaneously can be treated similarly by considering the whole sample path $y^i(t) \triangleq (y^i(1), y^i(2), \dots, y^i(t))$ at any time t .

Under assumptions (A1)-(A4) we study the following problems:

- (1) Will the estimates of the two agents eventually converge?
- (2) If they converge will they agree?

The model proposed in this report is similar to that of [3]. The only difference between the two models lies in (A1). In [3] all agents are Bayesian with the same view of the world; in the model proposed in this report the agents are Bayesian but have different views of the world (different P_i 's). Moreover each agent supposes that the other agent's a priori distribution is the same as his. Thus, the agents are unaware that they have different views of the world.* Because of (A1) the same message has different meanings for its sender and receiver: each agent interprets his data in terms of his own model and generates the conditional expectation according to his prior distribution P_i . The interpretation of the data and the message generation by each agent are considered in the next section.

*The case where the agents are aware that they have different views of the world reduces to the model of [3]; the agents can first agree on an a priori distribution P by negotiation (according to the method of [6]) and then proceed to solve the estimation problem.

3. DATA INTERPRETATION - MESSAGE GENERATION

In making his estimate \hat{x}_k^1 , agent 1 undertakes two logically distinct operations. First, he interprets his current data (his original observation and the received messages) in some consistent manner. This interpretation converts raw data into structured information. Secondly, based on the interpreted data an agent generates an estimate \hat{x}_k^1 of x , which he transmits to the other agent. Thus, a sequence of messages $\hat{x}_1^1, \hat{x}_1^2, \hat{x}_2^1, \hat{x}_2^2, \hat{x}_3^1, \hat{x}_3^2, \dots$ is generated as follows:

At $t=1$, 1's estimate is

$$\hat{x}_1^1 = E^1(x|Y^1) \quad p^1 \text{ a.s.} \quad (3-1)$$

(where E^1 denotes the expectation with respect to the probability measure p^1 induced by the distribution P^1 on Ω). The message \hat{x}_1^1 is transmitted to agent 2. Agent 2 interprets the message \hat{x}_1^1 as the realization of the random variable

$$\tilde{x}_1^1 = E^2(x|Y^1) \quad p^2 \text{ a.s.} \quad (3-2)$$

That is, according to agent 2 the realization \tilde{y}^1 of Y^1 is such that

$$\hat{x}_1^1 = \tilde{x}_1^1 = E^2(x|\tilde{y}^1) \quad p^2 \text{ a.s.} \quad (3-3)$$

At $t=2$, 2's estimate is

$$\hat{x}_1^2 = E^2(x|Y^2, \tilde{x}_1^1) \quad p^2 \text{ a.s.} \quad (3-4)$$

and this estimate is sent to 1 who interprets it as the realization of the random variable

$$\hat{x}_1^2 = E^1(x|Y^2, \hat{x}_1^1) \quad p^1 \text{ a.s.} \quad (3-5)$$

In other words, according to 1 the data \hat{y}^2 (the realization of Y^2) and \hat{x}_1^1 is such that

$$\hat{x}_1^2 = \hat{x}_1^2 = E^1(x|\hat{y}^2, \hat{x}_1^1) \quad p^1 \text{ a.s.} \quad (3-6)$$

In general, when agent 1 receives a message \hat{x}_k^2 he interprets it as

$$\hat{x}_k^2 = \hat{x}_k^2 = E^1(x|\hat{y}^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_k^1) \quad p^1 \text{ a.s.} \quad (3-7)$$

and then he generates

$$\hat{x}_{k+1}^1 = E^1(x|\hat{y}^1, \hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_k^2) \quad p^1 \text{ a.s.} \quad (3-8)$$

which he transmits to agent 2.

Agent 2 interprets this message as

$$\hat{x}_{k+1}^1 = \hat{x}_{k+1}^1 = E^2(x|\hat{y}^1, \hat{x}_1^2, \hat{x}_2^2, \hat{x}_3^2, \dots, \hat{x}_k^2) \quad p^2 \text{ a.s.} \quad (3-9)$$

and then transmits

$$\hat{x}_{k+1}^2 = E^2(x|\hat{y}^2, \hat{x}_1^1, \dots, \hat{x}_{k+1}^1) \quad p^2 \text{ a.s.} \quad (3-10)$$

to agent 1.

Equations 3-1 through 3-10 describe the rules according to which the sequence of messages $\hat{x}_1^1, \hat{x}_1^2, \hat{x}_2^1, \hat{x}_2^2, \dots$ is generated. These rules are well-defined because of assumption A2.

4. CONVERGENCE AGREEMENT

In this section we prove that under assumptions A1 through A4 and the rules by which messages are generated, one of two possible events occurs:

Either

1. The agents' estimates converge and the agents agree asymptotically,

or

2. The agents realize that they have different models and stop communicating any further.

To prove the result we re-examine the "data interpretation" which is one of the two distinct operations that each agent undertakes in making his estimate.

At first we let \underline{r} denote the realization of the random variable r .

When agent 2 receives message \hat{x}_k^1 he figures that Y^1 has realized a value \hat{y}_1 such that

$$E^2(x|\hat{y}_1, \hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_{i-1}^2) = \hat{x}_1^1 = E^1(x|\hat{y}_1, \hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_{i-1}^2) \quad \forall i < K \quad (4-1)$$

Similarly, when agent 1 receives message \hat{x}_k^2 he figures that Y^2 has realized \hat{y}^2 such that

$$E^1(x|\hat{y}_2^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_i^1) = \hat{x}_1^2 = E^2(x|\hat{y}_2^2, \hat{x}_1^1, \dots, \hat{x}_i^1) \quad \forall i < K \quad (4-2)$$

For each observation value $\underline{Y} \triangleq (Y_1, Y_2)$ let $\underline{Y}_1^1(\underline{Y})$ denote the set of all \tilde{Y}_1^1 that satisfy Eq. 4-1) and let $\underline{Y}_1^2(\underline{Y})$ be the set of all \tilde{Y}_1^2 that satisfy Eq. 4-2. Obviously these sets cannot increase, i.e.,

$$\underline{Y}_{k+1}^i(\underline{Y}) \subset \underline{Y}_k^i(\underline{Y}) \quad \forall k, i=1,2 \quad (4-3)$$

Thus, as k increases one of two cases can occur: Either

Case 1

$$\underline{Y}_1^i(\underline{Y}) \neq \emptyset \quad \forall i, \forall i \quad (4-4)$$

or

Case 2

At some step i

$$\underline{Y}_1^i(\underline{Y}) = \emptyset, \quad i=1 \text{ or } 2 \quad (4-5)$$

We shall analyze each case separately.

Case 1

When Eq. 4-4 is true we shall prove that the agents' estimates converge and the agents agree asymptotically.

To prove this result we first define the following σ -fields:

$$G_1^1 \triangleq \sigma^1(\tilde{x}_1^1, \tilde{x}_2^1, \dots, \tilde{x}_1^1, \tilde{x}_1^2, \tilde{x}_2^2, \dots, \tilde{x}_{1-1}^2) \quad (4-6)$$

$$F_i^2 \triangleq \sigma^1(Y^1, \tilde{x}_1^2, \tilde{x}_2^2, \dots, \tilde{x}_{1-1}^2) \quad (4-7)$$

$$\hat{G}_1^2 \triangleq \sigma^1 (\hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^1, \hat{x}_2^2, \hat{x}_2^2, \dots, x_1^2) \quad (4-8)$$

$$\hat{F}_1^2 \triangleq \sigma^1 (y^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^1) \quad (4-9)$$

$$\hat{G}_1^2 \triangleq \sigma^2 (\hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_1^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^1) \quad (4-10)$$

$$\hat{F}_1^2 \triangleq \sigma^2 (y^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^1) \quad (4-11)$$

$$\hat{G}_1^1 \triangleq \sigma^2 (\hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_{i-1}^2, \hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^1) \quad (4-12)$$

$$\hat{F}_1^1 \triangleq \sigma^2 (y^1, \hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_{i-1}^2) \quad (4-13)$$

$$\hat{G}_-^1 = \hat{G}_-^2 = \sigma^1 (\hat{x}_1^1, \hat{x}_2^1, \dots, \hat{x}_1^2, \hat{x}_2^2, \dots) \quad (4-14)$$

$$\hat{G}_-^2 = \hat{G}_-^1 = \sigma^2 (\hat{x}_1^2, \hat{x}_2^2, \dots, \hat{x}_1^1, \hat{x}_2^1, \dots) \quad (4-15)$$

The interpretation of $\hat{G}_1^1, \hat{G}_1^2, \hat{F}_1^1, \hat{F}_1^2$ is straightforward. The σ -field $\hat{G}_1^2 (\hat{G}_1^1)$ represents the view of agent 1(2) about the information available to agent 2(1) due to the messages generated and exchanged up until time $2i(2i-1)$. The σ -fields \hat{F}_1^1, \hat{F}_1^2 have a similar interpretation. We can now prove the following results:

Lemma 4-1

$$\{\hat{x}_1^1, \hat{G}_1^1, p^1\}, \{\hat{x}_1^2, \hat{G}_1^2, p^1\}, \{\hat{x}_1^2, \hat{G}_1^2, p^2\}, \{\hat{x}_1^1, \hat{G}_1^1, p^2\}$$

are uniformly interable martingales.

Proof

By definition

$$\tilde{x}_1^j = E^j \{x | F_1^j\} \quad p^j \text{ a.s.} \quad \forall j, j=1,2 \quad (4-16)$$

and

$$G_1^j \subset F_1^j \quad \forall j, j=1,2 \quad (4-17)$$

Because of Eqs. 3-11 and 3-12 and Theorem I.8.1 of [7] it follows that

$$\tilde{x}_1^j = E^j \{x | G_1^j\} \quad p^j \text{ a.s.} \quad \forall j, j=1,2 \quad (4-18)$$

The proof that $\{\tilde{x}_1^2, \tilde{G}_1^2, p^1\}, \{\tilde{x}_1^1, \tilde{G}_1^1, p^2\}$ are uniformly integrable martingales is similar. □

Theorem 4.1

The estimates $\tilde{x}_1^1, \tilde{x}_1^2, (\tilde{x}_1^2, \tilde{x}_1^1)$ of agent 1(2) converge $p^1 (p^2)$ a.s.

Proof

Follows from lemma 4.1 and the martingale convergence Theorem VII.4.1 of [7]. □

Let us denote by

$$\tilde{x}_\infty^1 = E^1 \{x | G_\infty^1\} \quad p^1 \text{ a.s.} \quad (4-19)$$

$$\tilde{x}_\infty^2 = E^1 \{x | G_\infty^2\} \quad p^1 \text{ a.s.} \quad (4-20)$$

$$\tilde{x}_\infty^2 = E^2 \{x | G_\infty^2\} \quad p^2 \text{ a.s.} \quad (4-21)$$

$$\tilde{x}_\infty^2 = E^2 \{x | G_\infty^1\} \quad p^2 \text{ a.s.} \quad (4-22)$$

the limits of the estimates of the two agents. Then, we have the following result.

Lemma 4.2

$$p^1 \{ \hat{x}_\infty^1 = \tilde{x}_\infty^2 \} = 1 \text{ and } p^2 \{ \hat{x}_\infty^2 = \tilde{x}_\infty^1 \} = 1 \quad (4-23)$$

Proof

Equation 4-23 follows from Eqs. 4-19 through 4-22 and 4-14 and 4-15. \square

So far we have shown that each agent's estimates converge and moreover each agent's estimate coincides asymptotically with his interpretation of the other agent's estimate within the terms of his own view of the world. The next theorem shows that asymptotically the two agents agree.

Theorem 4.2

Under Assumptions A1 through A4 and 4-4, the estimates of the two agents agree asymptotically.

Proof

When Eq. 4-4 is true, Eqs. 4-1 and 4-2 hold for all k , consequently

$$\left. \begin{array}{l} \hat{x}_\infty^1 = \tilde{x}_\infty^1 \\ \text{and} \\ \hat{x}_\infty^2 = \tilde{x}_\infty^2 \end{array} \right\} \quad (4-24)$$

Agreement then follows from Eq. 4-24 and Lemma 4.2 \square

The investigation of Case 1 is now complete.

Case 2

The following result is true in this case.

Theorem 4.3

If Eq. 4-5 is true the agents realize that they have different models and stop communicating any further.

Proof

If Eq. 4-5 occurs for agent i at some time \bar{t} , then agent i must assume that the sequence $\{\underline{y}_j^i(\underline{Y})\}_{j=1}^{\bar{t}}$ is "impossible" (i.e., an event of zero probability), or more reasonably agent i must assume that the two models p^1 and p^2 are different. Thus, further communication is not necessary (unless the agent is willing to modify his model and reinterpret the sequence of received messages).

Remark

If after $\underline{y}_i^1(\underline{Y}) = \phi$ agent i is willing to modify his model (his prior probability), then it can be shown, using the results of [6] and [3], that eventually asymptotic agreement can be achieved. The situation where agents are willing to modify their beliefs, after they receive an "impossible" sequence of messages according to their initial view of the world, is similar to that considered by Kreps and Wilson [8] who study dynamic games of perfect recall and determine "sequential equilibria" for these games. In Kreps and Wilson [8] the agents modify their beliefs (expressed by a behavioral strategy and a probability measure on the elements of every information set) whenever an information set of measure zero is reached in the game.

The investigation of Case 2 is now complete.

The results obtained so far show that the communicating agents either agree or realize that they have different models and terminate their communication. The example that follows shows that agreement may depend on the order of communication.

Example 4.1

Let $\Omega = [0,2] \times [0,3]$ (Fig. 4-1). Assume that P^1 is the Lebesgue measure (normalized to give $P(\Omega)=1$), and P^2 is such that

$$P^2(C_1) = 2/12, P^2(C_2) = 3/12, P^2(C_3) = 7/12.$$

The distribution on each of C_1, C_2, C_3 is uniform. Agents 1 and 2 try to estimate $I(A)$ where A is shown in Figure 4-1. At $t=0$, agent 1 observes $\{I(B_1), I(B_2)\}$ and agent 2 observes $\{I(C_1), I(C_2), I(C_3)\}$. Let $\omega \in B_1 \cap C_3$.




Figure 4-1. Parameters of Example 4.1

Consider first the situation where agent 1 sends his estimate \hat{x}_t^1 to agent 2 at times $t=1,3,5,\dots$, and agent 2 sends his estimate \hat{x}_t^2 to agent 1 at times $t=2,4,6,\dots$.

Then $\hat{x}_1^1(\omega) = E^1\{I(A)|I(B_1)\} = 1/2$. Agent 2 receives $1/2$ and interprets it as follows. He believes agent 1 observed $I(B_2)$ because only then

$$E^2\{I(A)|I(B_2)\} = 1/2 = \hat{x}_1^1(\omega).$$

Agent 2 initially observes $1(C_3)$. Consequently after he receives $\hat{x}_1^1(\omega)$ his new estimate is

$$\hat{x}_1^2(\omega) = E^2\{1(A)|1(B_2 \cap C_3)\} = 3/4 .$$

This estimate is transmitted to agent 1 who interprets it as follows: He concludes that agent 2 observed $1(C_2)$ because only then $E^1\{1(A)|1(B_1 \cap C_2)\} = 3/4$. Then, agent 1's estimate is

$$\hat{x}_2^1(\omega) = E^1\{1(A)|1(B_1 \cap C_2)\} = 3/4$$

and this estimate is transmitted to agent 2. From that point on further communication does not convey any more information, so the agents' estimates agree and

$$\hat{x}^1 = \hat{x}^2 = 3/4 .$$

Consider now the situation where agent 2 sends his estimate \hat{x}_t^2 to agent 1 at times $t=1,3,5,\dots$ and agent 1 sends his estimate \hat{x}_t^1 to agent 2 at times $t=2,4,6,\dots$. Again, initially agent 1 observes $1(B_1)$ and agent 2 observes $1(C_3)$. Hence

$$\hat{x}_1^2(\omega) = E^2\{1(A)|1(C_3)\} = 1/2 .$$

Agent 2 receives $\hat{x}_1^2(\omega)$ and interprets it as follows: He believes that agent 2 observed either $1(C_2)$ or $1(C_3)$ because only then

$$E^1\{1(A)|1(B_1 \cap C_2)\} = E^1\{1(A)|1(B_1 \cap C_3)\} = 1/2 .$$

Then agent 1 forms the new estimate

$$\hat{x}_1^1(\omega) = \frac{1}{2} E^1\{1(A)|1(B_1 \cap C_2)\} + \frac{1}{2} E^1\{1(A)|1(B_1 \cap C_3)\} = \frac{1}{2}$$

which he communicates to agent 2.

Agent 2 can not interpret consistently this message for the following reason:

He knows that agent 1 observed either $1(B_1)$ or $1(B_2)$. He also knows that agent 1 knows that agent 2 observed either $1(C_2)$ or $1(C_3)$. Hence if agent 1 observed $1(B_1)$ he should transmit to agent 2 the message $\hat{x}_1^1(\omega) = 0.4$; on the other hand if agent 1 observed $1(B_2)$ he should transmit to agent 2 the message $\hat{x}_1^1(\omega) = 0.6$. The message $\hat{x}_1^1(\omega) = 0.5$ has a zero probability according to agent 2's view of the world, consequently, agent 2 realizes that he does not have the same model as agent 1 and any further communication is not necessary. \square

Example 4.1 shows that it is not possible to partition the space Y of observations $\underline{Y} \triangleq (Y^1, Y^2)$ into sets Y^C, Y^D such that

$$Y^C \cup Y^D = Y$$

$$Y^C \cap Y^D = \emptyset$$

and so that if $\underline{Y} \in Y^C$ agreement is always achieved whereas if $\underline{Y} \in Y^D$ the agents realize eventually that they have different models.

Finally, we restrict attention to the Gaussian case. By this we mean that (i) the measurements are jointly Gaussian, (ii) Assumptions A1 through A4 hold. Then we have the following result:

Lemma 4.3

For the Gaussian case agreement is always achieved when the order of communication is fixed. However, when the two agents communicate simultaneously, then agreement occurs with zero probability.

Proof

When the order of communication is fixed agreement follows from Lemma 3 of [3] and the fact that with only one measurement and two agents having the same view of the world it takes only one communication to reconstruct the centralized estimate.

When the agents communicate simultaneously then, in view of the results of [3], after one communication each agent believes that he has reconstructed the centralized estimate (according to his own model), thus the two agents do not communicate any further and agreement occurs with zero probability.

5. CONCLUSIONS

This report extended the results of [3] to the case where agents with different beliefs (different prior probabilities) exchange information.

The process of expectation formation for such systems was described, and it was shown that agreement depends upon the order of communication between the agents.

REFERENCES

1. Aumann, R.J., "Agreeing to Disagree," Annals of Statistics, Vol. 4, No. 6, 1976, pp. 1236-1239.
2. Geanakopoulou, J.D. and H.M. Polemarchakis, "We Can't Disagree Forever," Technical Report No. 277, Institute for Mathematical Studies in the Social Sciences, Stanford University, December 1978.
3. Borkar, V. and P. Varaiya, "Asymptotic Agreement in Distributed Estimation", IEEE Trans. on Automatic Control, Vol. AC-27, No. 3.
4. Tsitsiklis, J.N. and M. Athans, "Convergence and Asymptotic Agreement in Distributed Decision Problems," LIDS-P-1185, MIT Laboratory for Information and Decision Systems, Cambridge, Massachusetts, March 1982.
5. Washburn, R.B. and D. Teneketzis, "Asymptotic Agreement Among Communicating Decisionmakers," ALPHATECH Technical Report TR-145, Burlington, Massachusetts, February 1983. Also submitted to Mathematics of Operations Research.
6. DeGroot, M.H., "Reaching a Consensus," Journal of the American Statistical Association, March 1974, Volume 69, No. 345.
7. Doob, J.L., Stochastic Processes, J. Wiley, 1953.
8. Kreps, D.M. and R. Wilson, "Sequential Equilibria," Technical Report No. 340, Institute for Mathematical Studies in the Social Sciences, Stanford University, December 1980.

APPENDIX C

CONSENSUS IN DISTRIBUTED ESTIMATION*

*This research was also supported in part by ONR contract N00014-80-C-0507, N00014-82-C-0693, and JSEP contract F49620-79-C-0178.

C-0

Consensus in Distributed Estimation¹

Demosithenis Teneketzis

Department of Electrical Engineering
and Computer Science
University of Michigan
Ann Arbor, MI 48109

Pravin Varaiya

Department of Electrical Engineering
and Computer Sciences
and Electronics Research Laboratory
University of California
Berkeley, CA 94720

ABSTRACT

A team must agree on a common decision to minimize the expected cost. Different team members have different observations relating to the 'state of the world', and they may also have different prior beliefs. To reach a consensus they exchange tentative decisions based on their current information. Two questions are discussed: When do the individual estimates converge? If they converge, will a consensus be reached?

PAGES 3,5,6,8,11,32
ARE
MISSING
IN
ORIGINAL
DOCUMENT

1. INTRODUCTION

A team or committee of N people, indexed $i = 1, \dots, N$, must agree on a common decision u to be selected from a pre-specified set U so as to minimize the cost

$$J(\omega, u) \quad (1)$$

where J is a real valued function of the 'state of the world' $\omega \in \Omega$, and the decision u . Initially, different people have different information relating to ω . This is modeled by stipulating that person i observes the value of the random variable $Y_i = Y_i(\omega)$. Everyone knows that i knows Y_i , although $j, j \neq i$, does not know what the value of Y_i actually is. Everyone knows the function J .

Each person has a prior belief concerning ω . We stipulate that i 's prior belief is summarized by the probability distribution P^i on (Ω, \mathcal{F}) where \mathcal{F} is the σ -field of events. If $P^1 = \dots = P^N$, we say that the beliefs are *consistent*; otherwise they are *inconsistent*.

Since initially different people have different information, and also because their beliefs may be inconsistent, their estimates of the best decision will also be different. To arrive at a consensus decision it is necessary for them to share information. We suppose that this information is shared by means of the following procedure.

Consider person i . In the first round he makes an estimate $u_i(1)$ which is based on his initial data Y_i , and he communicates this estimate to some or all of the other members. By the time i makes his second estimate, he will have received the estimates of some of the others. More generally, denote by $D_i(t-1)$ the messages received by i from the others before i makes his t^{th} estimate $u_i(t)$. That estimate will be based on Y_i and $D_i(t-1)$. We assume that i communicates all his estimates to a fixed set of the other people, and that there is a message transmission delay of one time unit.

Our aim is to discuss two questions: Will each person's estimate converge as $t \rightarrow \infty$? If the individual estimates converge, will they reach a common limit? To formulate these questions mathematically, we need to specify how each person estimates the best decision based on the data available to him. This is done in Section 2. Once this is done, it turns out that the answers depend crucially upon whether the prior beliefs are consistent or inconsistent. The consistent case is considered in Section 3, and the inconsistent case in Section 4. Section 5 outlines some directions for further research.

2. ESTIMATION SCHEMES

Several different estimation schemes have been considered in the literature.

Borkar and Varaiya [2] consider the situation where the committee wants to estimate a random variable X , and they suppose that the i^{th} estimate made by i , $u_i(t)$, is the conditional mean of X given the available data, i.e.,

$$u_i(t) = E^i \{X \mid Y_i, D_i(t-1)\}. \quad (2)$$

Here E^i denotes expectation with respect to P^i . We will see later that the right hand side of (2) has to be interpreted carefully when the beliefs are inconsistent. For the moment observe that the estimate given by (2) is also the decision that minimizes the (expected value of the) cost function

$$J(\omega, u) := |X(\omega) - u|^2$$

when the information available is $\{Y_i, D_i(t-1)\}$. Aumann [1], and Geanakoplos and Polemarchakis [4] consider the situation in which the group wants to estimate the probability that a particular event $F \in \underline{F}$ has occurred. This is a special case of (2) with $X = 1(F)$. The set Ω of all possible states is finite in [1] and [4]. Tsitsiklis and Athans [7] consider the situation described in the introduction. Sebenius and Geanakoplos [5]

$$u_i(t) = E\{X \mid \underline{Y}_i(t-1)\} . \quad (4)$$

Let $\underline{Y}_i(\infty) := \bigvee_t \underline{Y}_i(t)$. Since $\underline{Y}_i(t)$ is an increasing sequence, it follows from the martingale convergence theorem that

$$u_i(t) \rightarrow u_i(\infty) \text{ a.s.}; u_i(\infty) := E\{X \mid \underline{Y}_i(\infty)\} . \quad (5)$$

Thus the individual estimates do converge.

Next we investigate whether the limiting estimates agree. Suppose i communicates his estimates to j . Then $u_i(t)$ is $\underline{Y}_j(t+1)$ -measurable. From (5) it follows that $u_i(\infty)$ is $\underline{Y}_j(\infty)$ -measurable, and so,

$$u_i(\infty) = E\{u_i(\infty) \mid \underline{Y}_i(\infty) \cap \underline{Y}_j(\infty)\} . \quad (6)$$

Suppose there is a communication ring $i_1, \dots, i_n = i_1$. This is a not necessarily distinct sequence of persons such that i_k communicates his estimates to i_{k+1} . Then, according to (6), we must have

$$u_{i_k}(\infty) = E\{u_{i_{k+1}}(\infty) \mid \underline{Y}_{i_k}(\infty) \cap \underline{Y}_{i_{k+1}}(\infty)\} , \quad k = 1, \dots, n , \quad (7)$$

where $i_{n+1} := i_1$. It is quite easy to show [2, Lemma 2] that (7) implies

$$u_{i_1} = \dots = u_{i_n} .$$

so that the asymptotic estimates of the members of a communication ring agree. This suggests the main result of [2]:

Theorem 1.

If the estimates of i are given by (2), then each person's estimate converges. Moreover, if everyone in the team is a member of the same communication ring then the limiting estimates agree.

Proof:

Teneketzis [8] show also that the agreement condition for rings is satisfied by decision rules which are optimal in the sense defined below:

Proposition 1.

Suppose that d is a decision rule such that $\sigma(d(F')) \subset F'$ for all σ -fields $F' \subset F$. Then d satisfies the agreement condition for rings if and only if there is a partial ordering \leq of the set of functions $\{d(F') : F' \subset F\}$ such that $d(F')$ is the maximum element of $\{d(G) : G \subset F, \sigma(d(G)) \subset F'\}$ with respect to \leq .

Proof:

See Appendix C.

In many cases as in [2], [7] the partial order relation is defined in terms of a scalar cost function. The following proposition proves that decision rules defined by such cost functions satisfy the agreement condition for rings, provided that the decision includes a tie-breaking rule of the cost function has more than one minima.

Proposition 2.

Suppose that the decision functions take values in a set U . Let L be a real-valued functional of F -measurable decision functions $\delta : \Omega \rightarrow U$. For each F' let $D(F')$ be the set of F' -measurable decision functions δ such that $L(\delta) \leq L(\delta')$ for all measurable $\delta' : \Omega \rightarrow U$. Assume that U is partially ordered by \leq' and that for each F' , there is a $\delta \in D(F')$ such that $\delta' \in D(F')$ implies $\delta'(\omega) \leq' \delta(\omega)$ for all $\omega \in \Omega$. The decision rule $d(F')$ which assigns this $\delta \in D(F')$ to F' satisfies the agreement condition for rings.

Proof:

3.3. COMMON KNOWLEDGE

The main feature of the estimation schemes presented in [1], [2], [4], [8] is the following:

If all team members use the same decision rule, if everyone in the team is a member of the same communication ring and if common knowledge decisions agree, then all team members agree on the same decision. The common decision is the decision based on the ultimate common knowledge (common information) of the team members.

Thus, it appears appropriate to define common knowledge at this point, and to show that the definitions of common information given in [1], [2], [8], [9] are essentially equivalent and lead to the same results.

Aumann [1] represents information by a partition \underline{P} on the sample space Ω . Borkar-Varaiya [2] and Washburn-Teneketzis [8] represent information by σ -fields contained in \underline{F} . It can be shown that these two representations are essentially equivalent. The partition \underline{P} is a collection $\{E_1, E_2, \dots\}$ of mutually disjoint events whose union is the whole sample space. To a partition \underline{P} there corresponds a unique σ -field \underline{F} , namely the σ -field generated by the events in \underline{P} . Each $E_i \in \underline{P}$ is an atom of \underline{F} . If $\underline{P}_1 = \{E_1, E_2, \dots\}$ and $\underline{P}_2 = \{G_1, G_2, \dots\}$ then one can define a third partition \underline{P}_3 which is the finest partition contained in \underline{P}_1 and \underline{P}_2 and is denoted by $\underline{P}_1 \wedge \underline{P}_2$. If \underline{P}_1 and \underline{P}_2 correspond to the σ -fields \underline{F}_1 and \underline{F}_2 then $\underline{P}_1 \wedge \underline{P}_2$ corresponds to $\underline{F}_1 \wedge \underline{F}_2$. Aumann [1] defines an event E to be common knowledge to team members 1 and 2 (with information \underline{P}_1 and \underline{P}_2 respectively) at ω if there is an atom $\dot{G} \in \underline{P}_1 \wedge \underline{P}_2$ such that $\omega \in \dot{G} \subset E$. If \underline{F}_1 and \underline{F}_2 are the σ -fields corresponding to \underline{P}_1 and \underline{P}_2 respectively, then the definitions of common knowledge at ω given in [2], [8], namely that there is $\dot{G} \in \underline{F}_1 \wedge \underline{F}_2$ and $\omega \in \dot{G} \subset E$, are equivalent to Aumann's definition. Let us say that

the event E is common knowledge to the team members 1 and 2 if it is common knowledge at each $\omega \in E$. Then E is common knowledge to 1 and 2 if and only if it belongs to the σ -field generated by $\underline{P}_1 \wedge \underline{P}_2$, namely $\underline{F}_1 \wedge \underline{F}_2$.

Milgrom [9] characterizes common knowledge by

(i) associating with each event E another event K_E with the interpretation

$$K_E = \{ \omega \in \Omega : E \text{ is common knowledge at } \omega \}$$

and

(ii) considering the following four conditions:

$$(C1) K_E \subset E$$

$$(C2) \forall \omega \in K_E, \forall i, \text{ if } \omega \in F_i, F_i \subset K_E$$

$$(C3) E_1 \subset E_2 \implies K_{E_1} \subset K_{E_2}$$

$$(C4) [\forall i, \forall \omega \in E, \text{ if } \omega \in F_i, F_i \in E] \implies E = K_E.$$

Condition (C1) asserts that an event E is common knowledge only if it actually occurs. Condition (C2) implies that if E is common knowledge then every team member knows that E is common knowledge. Conditions (C1) and (C2) imply that E is common knowledge only if E occurs, each team member knows E , each knows that all know E and so on. Condition (C3) implies that wherever E_1 is common knowledge any logical consequence of E_1 is also common knowledge. Condition (C4) asserts that public events are common knowledge whenever they occur. A public event is defined by the antecedent in (C4); it is an event which if it occurs will be known to every team member. Milgrom ([9]) shows that his characterization of common knowledge is equivalent to Aumann's definition.

mation in a way that can be stated succinctly using common knowledge: a contingent allocation f is efficient if there is no other allocation ν such that it is common knowledge that all agents prefer f to ν . Milgrom ([9]-[10]) used the idea of common knowledge to analyze a rational expectations trading model. He showed that when traders exchange a risky security on the basis of private information then they "agree to disagree." (i.e., no trade takes place). Kreps, Milgrom, Roberts and Wilson [12] consider finite repetitions of the well-known prisoners' dilemma game. A common observation in experiments involving finite repetitions of the prisoners' dilemma is that players do not always play the single period dominant strategies but instead achieve some measure of cooperation. Kreps and his co-authors in [12] show that the lack of common knowledge about one or both players' options, motivation or behavior can explain the observed cooperation.

4. INCONSISTENT BELIEFS:

The analysis is quite different when the beliefs are inconsistent. The discussion in this section is initially based on Teneketzis and Varaiya [6]. Then the results of [6] are extended to the case of a general decision rule d . To keep the notation simple assume there are only two persons, *Alpha* and *Beta*. Initially, *Alpha* observes the random variable A and *Beta* observes B . Both wish to estimate the random variable X . We also assume that Ω is finite. The prior probabilities of *Alpha* and *Beta* are denoted P^a, P^b respectively.

For $t = 1, 2, \dots$ the t^{th} estimate by *Alpha* (*Beta*) is denoted α_t (β_t). α_t is the conditional expectation of X given the observations $A, \beta_1, \dots, \beta_{t-1}$. After α_t has been calculated it is communicated to *Beta* whose t^{th} estimate is the conditional expectation of X given $B, \alpha_1, \dots, \alpha_t$. Once β_t is evaluated it is communicated to *Alpha* who incorporates it into the estimate α_{t+1} , and the procedure is repeated.

To complete the specification we assume that the estimation procedures followed by *Alpha* and *Beta* are consistent with their own prior models. That is, each assumes the other's model to be the same as his own. Consider *Alpha*. When he receives *Beta's* estimate β_{L-1} , *Alpha* interprets it as if it were based on P^α rather than on P^β . Thus *Alpha* assumes that *Beta's* estimate is a realization of the random variable

$$\hat{\beta}_{L-1} := E^\alpha \{X \mid B, \alpha_1, \dots, \alpha_{L-1}\}.$$

Subsequently, *Alpha* calculates α_L ,

$$\alpha_L := E^\alpha \{X \mid A, \hat{\beta}_1, \dots, \hat{\beta}_{L-1}\}.$$

Symmetrically, *Beta* interprets α_L as

$$\hat{\alpha}_L := E^\beta \{X \mid A, \beta_1, \dots, \beta_{L-1}\},$$

and calculates β_L by

$$\beta_L := E^\beta \{X \mid B, \hat{\alpha}_1, \dots, \hat{\alpha}_L\}.$$

There is a more revealing description of the functional dependence of these estimates. Suppose a particular realization $\bar{\omega} = (\bar{A}, \bar{B})$ has occurred. Since *Alpha* observes \bar{A} , he concludes that $\bar{\omega} \in \Omega_1^\alpha := \{(A, B) \mid A = \bar{A}\}$ and so his first estimate equals

$$\bar{\alpha}_1 = E^\alpha \{X \mid A = \bar{A}\} = E^\alpha \{X \mid \omega \in \Omega_1^\alpha\}.$$

Alpha transmits the number $\bar{\alpha}_1$ to *Beta*. *Beta* interprets it as a realization of the random variable

$$\hat{\alpha}_1 = E^\beta \{X \mid A\},$$

and so he infers that $\bar{\omega} \in \Omega_1^\beta := \{\omega \mid \hat{\alpha}_1(\omega) = \bar{\alpha}_1, B = \bar{B}\}$, and his first estimate takes the value

$$\bar{\beta}_1 = E^\beta \{X \mid \omega \in \Omega_1^\beta\}.$$

This value is communicated to *Alpha*.

At the beginning of the t^{th} round, Alpha starts with the inference $\bar{\omega} \in \Omega_t^\alpha$ when he receives the estimate $\bar{\beta}_{t-1}$. He interprets it as a realization of the random variable

$$\hat{\beta}_{t-1} = E^\omega\{X \mid B, \alpha_1, \dots, \alpha_{t-1}\}$$

and so Alpha concludes that $\bar{\omega} \in \Omega_t^\alpha := \{\omega \mid \omega \in \Omega_{t-1}^\alpha, \hat{\beta}_{t-1}(\omega) = \bar{\beta}_{t-1}\}$. Hence Alpha's t^{th} estimate takes the value

$$\bar{\alpha}_t = E^\omega\{X \mid \omega \in \Omega_t^\alpha\}$$

which is communicated to Beta. Whereupon Beta interprets it as a realization of

$$\hat{\alpha}_t = E\{X \mid A, \beta_1, \dots, \beta_{t-1}\},$$

concludes that $\bar{\omega} \in \Omega_t^\beta := \{\omega \mid \omega \in \Omega_{t-1}^\beta, \hat{\alpha}_t(\omega) = \bar{\alpha}_t\}$ and evaluates his t^{th} estimate as

$$\bar{\beta}_t = E^\beta\{X \mid \omega \in \Omega_t^\beta\}.$$

Thus, as expected, the uncertainty diminishes with each exchange, $\Omega_{t+1}^\alpha \subset \Omega_t^\alpha$, $\Omega_{t+1}^\beta \subset \Omega_t^\beta$. From the description above we also see that if for some k either $\Omega_{k+1}^\alpha = \Omega_k^\alpha$ or $\Omega_{k+1}^\beta = \Omega_k^\beta$, then $\Omega_t^\alpha = \Omega_{k+1}^\alpha$ and $\Omega_t^\beta = \Omega_{k+1}^\beta$ for $t > k+1$. Hence for $t > T$ (which cannot exceed the number of distinct elements in Ω), Ω_t^α and Ω_t^β become constant. These limit sets depend upon the realization ω . Call them $\Omega^\alpha(\omega)$ and $\Omega^\beta(\omega)$ respectively.

There are two possibilities. The first is that $\Omega^\alpha(\omega) = \phi$ and $\Omega^\beta(\omega) = \phi$. This happens because at some stage the message $\bar{\beta}_{t-1}$ received by Alpha is "impossible:" there is no $\bar{\omega}$ such that $\hat{\beta}_{t-1}(\bar{\omega}) = \bar{\beta}_{t-1}$; or the message $\bar{\alpha}_t$ received by Beta is "impossible:" there is no $\bar{\omega}$ such that $\hat{\alpha}_t(\bar{\omega}) = \bar{\alpha}_t$. Alpha and Beta must realize that their prior models are inconsistent. Let Ω_I be the set of all realizations that lead to this outcome.

The second possibility is that $\Omega^\alpha(\omega) \neq \phi$ and $\Omega^\beta(\omega) \neq \phi$. In this case for $t > T$ the estimates stop changing: $\hat{\beta}_t(\omega) = \hat{\beta}_\bullet(\omega)$, $\alpha_t(\omega) = \alpha_\bullet(\omega)$, $\hat{\alpha}_t(\omega) = \hat{\alpha}_\bullet(\omega)$, $\beta_t(\omega) = \beta_\bullet(\omega)$. Since for every t , $\hat{\beta}_t(\omega) = \beta_t(\omega)$ and $\hat{\alpha}_t(\omega) = \alpha_t(\omega)$, it follows that

$$\hat{\beta}_t(\omega) = \beta_t(\omega), \hat{\alpha}_t(\omega) = \alpha_t(\omega).$$

On the other hand, since $\hat{\beta}_t$ and α_t are based on the same model, namely P^α , it follows from Theorem 1 that $\hat{\beta}_t(\omega) = \alpha_t(\omega)$. For the same reason $\hat{\alpha}_t(\omega) = \beta_t(\omega)$. Thus if $\omega \in \Omega_{II} := \Omega - \Omega_I$, there is agreement $\alpha_t(\omega) = \beta_t(\omega)$ for $t > T$. It is worth emphasizing that this agreement need not be a reflection of the consistency of the two models P^α , P^β . Rather agreement occurs because within each person's model there is sufficient "uncertainty" to permit the reconciliation of the other's messages with his own observation. One might say that agreement could result from two wrong arguments. We summarize the preceding analysis as follows:

Theorem 3

The set of events Ω decomposes into two disjoint subsets Ω_I and Ω_{II} . After T exchanges, if $\omega \in \Omega_I$ both agents realize their models are inconsistent, whereas if $\omega \in \Omega_{II}$ the two estimates coincide.

The result is fragile. In particular, whether a realization ω ends in agreement or in impasse can depend upon the order of communication between Alpha and Beta as demonstrated by the following example:

Example

Take $\Omega = [0,2] \times [0,3]$, suppose Alpha observes

$$A = \{1(a_1), 1(a_2)\}$$

and Beta observes

$$B = \{1(b_1), 1(b_2), 1(b_3)\}$$

and suppose X is the indicator function of the shaded region as shown in figure 1. Assume that ω is uniformly distributed under P^α , whereas under P^β .

$$P^{\theta}(b_1) = 2/12, P^{\theta}(b_2) = 3/12, P^{\theta}(b_3) = 7/12$$

and within each b_i , ω is uniformly distributed. Suppose $\bar{\omega} \in a_1 \cap b_3$ and that Alpha communicates first. Then

$$\bar{a}_1 = E(X|\omega \in a_1) = 1/2$$

Beta interprets this as a realization of

$$\hat{a}_1 = E(X|I(a_1)I(a_2))$$

Since $E^{\theta}(x|\omega \in a_1) = 5/12$, $E^{\theta}(x|\omega \in a_2) = 1/2$, upon learning that $\bar{a}_1 = 1/2$, Beta concludes that $\bar{\omega} \in a_2$, and since he has observed that $\bar{\omega} \in b_3$ his estimate is

$$\bar{\beta}_1 = E^{\theta}(X|\omega \in a_2 \cap b_3) = 3/4$$

Alpha interprets $\bar{\beta}_1$ as a realization of $E^*(x|\omega \in a_1, B)$. Since

$$E^*(x|\omega \in a_1 \cap b_1) = 1/2,$$

$$E^*(x|\omega \in a_1 \cap b_2) = 3/4,$$

$$E^*(x|\omega \in a_1 \cap b_3) = 1/2,$$

Alpha concludes that $\bar{\omega} \in a_1 \cap b_2$, hence

$$\bar{a}_2 = E^*(X|\omega \in a_1 \cap b_2) = 3/4.$$

Evidently, $\bar{\beta}_2 = \bar{\beta}_3 = \dots = \bar{a}_2 = \bar{a}_3 = \dots = 3/4$ and there is agreement. (Note that Alpha believes that $\omega \in a_1 \cap b_2$, Beta believes that $\bar{\omega} \in a_2 \cap b_3$, in fact $\bar{\omega} \in a_1 \cap b_3$.)

Now suppose again that $\bar{\omega} \in a_1 \cap b_3$, but this time Beta communicates first. Then

$$\bar{\beta}_1 = E^{\theta}(X|\omega \in b_3) = 1/2$$

Since

$$E^*(X|\omega \in b_1) = 1/4,$$

$$E^*(X|\omega \in b_2) = E^*(X|\omega \in b_3) = 1/2,$$

upon learning $\bar{\beta}_1 = 1/2$, Alpha concludes that $\bar{\omega} \in b_2 \cup b_3$, then his estimate is

$$\bar{a}_1 = E^*(X|\omega \in a_1 \cap (b_2 \cup b_3)) = 1/2.$$

But Beta expects \bar{a}_1 to take on the value

$$E^{\beta}(X|\omega \in a_1 \cap (b_2 \cup b_3)) = 0.4$$

or

$$E^{\beta}(X|\omega \in a_1 \cap (b_2 \cup b_3)) = 0.6 .$$

Thus *Beta* concludes that the models are inconsistent.

The results of Teneketzis and Varaiya [6] can be extended to the case where the decision rule is a general function d as in section 3.2. We discuss this case next.

Assume the same model as in Teneketzis and Varaiya [6] and suppose the estimates a_i and β_i are generated by the decision rule d given the observations $A, \beta_1, \beta_2, \dots, \beta_{L-1}$ and $B, a_1, a_2, \dots, a_{L-1}$ respectively. Suppose the decision rule d satisfies the agreement condition:

$$\begin{aligned} &\text{for all } \underline{G}_1, \underline{G}_2 \subset \underline{F}, \\ &\sigma(d(\underline{G}_2)) \subset \underline{G}_1 \subset \underline{G}_2 \implies d(\underline{G}_1) = d(\underline{G}_2). \end{aligned} \quad (15)$$

Under the assumptions above one can prove the following result:

Theorem 4.

If Ω is finite and the decision rule d satisfies the agreement condition (15), then either the estimates a and β agree after a finite number of communications or *Alpha* and *Beta* realize that their models are inconsistent.

Proof:

See Appendix F.

As pointed out in the discussion previously the investigation of convergence and agreement of the estimates can proceed in two steps:

1. Determine what each team member's model predicts about the evolution and the outcome of the estimation process.

2. Examine how these predictions compare with what actually happens during the estimation process.

For finite Ω , the result of theorem 4 is true for rules that obey the agreement condition for a very simple reason. If a team member's view of the world is consistent with reality, then agreement must result after a finite number of communications because this is what is predicted by the team member's model; anything else would be inconsistent.

5. CONCLUDING REMARKS

Recall the discussion in Section 3.1 and 3.2. There a consensus is reached via a sequence of exchanges of tentative decisions. The information available to a person increases with each message exchange and the limiting consensus decision is based on the information common to all in the sense that $d_1(\infty) = \dots = d_N(\infty)$ is measurable with respect to $\underline{Y}_1(\infty) \cap \dots \cap \underline{Y}_N(\infty)$. A consensus can also be reached if all people share their initial private data $\underline{Y}_1, \dots, \underline{Y}_N$. We may call this consensus the *full information* decision. It turns out that the consensus reached by exchanging tentative decisions need not coincide with the full information decision. However, within a rather simple model, Geanakoplos and Polemarchakis [4] have shown that the two decisions are "almost always" the same. It would be worth investigating this in a more general setting.

Secondly, even when the two decisions are the same, it does not follow that all people obtain the full information, i.e., it need not be the case that $\underline{Y}_1(\infty) = \sigma\{\underline{Y}_1, \dots, \underline{Y}_N\}$. If $\underline{Y}_1(\infty)$ is a proper subset of $\sigma\{\underline{Y}_1, \dots, \underline{Y}_N\}$, then one could argue that reaching consensus via exchange of tentative decisions requires a transfer of less information than the exchange of all private information. This too is worth further investigation.

Recall now the discussion dealing with the case of inconsistent beliefs. The most interesting finding is that *Alpha* and *Beta* can exchange statements about X and eventually agree even when their views are different. Thus paradoxically, the realization that these views are different is only reached when further communication becomes impossible. This raises several basic and knotty issues that need further investigation.

One can readily imagine situations where the most important thing is to determine whether or not the beliefs are inconsistent. In the communication setup of Section 4 the realization that beliefs are inconsistent is fortuitous—it happens only if *Alpha* and *Beta* reach an impasse. How should one structure the set of message exchanges so as to expedite the reaching of an impasse?

Suppose now that *Alpha* and *Beta* do reach an impasse ($\omega \in \Omega_I$). Our analysis stops at this point, but there are two directions that can be pursued. First, observe that with the realization that their beliefs are different comes the understanding that they have "misread" each other's messages (i.e., they now know that $\hat{\beta}_t \neq \beta_t$ and $\hat{\alpha}_t \neq \alpha_t$), and consequently their estimates have been "biased." To eliminate this bias each needs to learn what the other's view is. A straightforward way of permitting such learning is to suppose that from the beginning *Alpha* admits that *Beta*'s model P^β might be any one of a known set \underline{P}^β of models and there is a prior distribution on \underline{P}^β reflecting *Alpha*'s initial judgment about *Beta*'s model; a symmetrical structure is formulated for *Beta*. Within such a framework it seems reasonable to conjecture that each agent will correctly read the other's message and his sequence of estimates will converge. But if their models are different then the limiting estimates may differ, and a consensus will not emerge.

Suppose, however, that *Alpha* and *Beta* want to reach a consensus. To reach a consensus one or both must change their models. One can imagine many different ways in

which this can be done. For example, De Groot [3] proposes that each person tells the others what his prior probability is, and he proposes an *ad hoc* behavioral rule whereby each person adjusts his model to a weighted average of the others' models. This is not very satisfactory in situations where communicating one's prior beliefs is not practicable.

REFERENCES

- [1] Aumann, R.J., "Agreeing to disagree," *Annals of Statistics*, Vol.4(6), 1976, pp.1236-1239.
- [2] Borkar, V. and P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Trans on Automatic Control*, Vol.AC-27(3), 1982, pp.650-655.
- [3] De Groot, M.H., "Reaching a consensus," *J. American Statistical Association*, Vol.69, No.345, March 1974, pp.118-121.
- [4] Geanakoplos, J.D. and H.M. Polemarchakis, "We can't disagree forever," *Journal of Economic Theory*, Vol.26, 1982, pp.192-200.
- [5] Sebenius, J.K. and J. Geanakoplos, "Don't bet on it: Contingent agreements with asymmetric information," *J. American Statistical Association*, Vol.78, No.382, June 1983, pp.424-236.
- [6] Teneketzis, D. and P. Varaiya, "Consensus in distributed estimation with inconsistent beliefs," *Systems and Control Letters*, 4, 1984, pp.217-221.
- [7] Tsitsiklis, J.N. and M. Athans, "Convergence and asymptotic agreement in distributed decision problems," *IEEE Trans on Automatic Control*, Vol.AC-29(1), 1984, pp.42-50.

- [8] Washburn, R.B. and D. Teneketzis, "Asymptotic agreement among communicating decisionmakers," *Stochastics*, Vol.13, 1984, pp.103-129.
- [9] Milgrom, P., "An axiomatic characterization of common knowledge," *Econometrica*, Vol.49, No.1, January 1981.
- [10] Milgrom, P. and N. Stokey, "Information, trade and common knowledge," Center for Math. Studies, Discussion paper No.377R, Northwestern University, November 1979.
- [11] Wilson, R., "Information, efficiency and the core of an economy," *Econometrica*, 45 (1978), pp.807-816.
- [12] Kreps, D.M., P. Milgrom, J. Roberts and R. Wilson, "Rational cooperation in the finitely-repeated prisoners' dilemma," *Journal of Economic Theory*, 27, (1982) pp.245-252.

APPENDIX A

Proof of Theorem 1

Convergence of each member's estimates follows from the Martingale Convergence Theorem. The proof of the rest of the Theorem proceeds in several steps. Consider two agents i and j and let $G_i(t)$ denote the σ -field generated by the transmission and reception of messages from agent i up to time t . That is,

$$G_i(t) = \sigma\{u_i(1), \dots, u_i(t-1), \dots, u_{i-1}(1), \dots, u_{i-1}(t-1), u_i(1), \dots, u_i(t), u_{i+1}(1), \dots, u_{i+1}(t-1), \dots, u_q(1), \dots, u_q(t-1)\}$$

$$G_j(t) = \sigma\{u_i(1), \dots, u_i(t-1), \dots, u_{j-1}(1), \dots, u_{j-1}(t-1), u_j(1), \dots, u_j(t), u_{j+1}(1), \dots, u_{j+1}(t-1), \dots, u_q(1), \dots, u_q(t-1)\}$$

Define S^i to be the event that agent i sends messages to j infinitely often. Then

Lemma A1

Both $u_i(\infty) \mathbb{I}(S^i)$ and $u_j(\infty) \mathbb{I}(S^j)$ are common knowledge for $G_i(\infty)$ and $G_j(\infty)$. Moreover,

$$u_i(\infty) \mathbb{I}(S^i) = E(X | G_i(\infty) \cap G_j(\infty)) \mathbb{I}(S^i) \quad \text{a.s.} \quad (\text{A.1})$$

and

$$u_i(\infty) \mathbb{I}(S) = u_j(\infty) \mathbb{I}(S) \quad (\text{A.2})$$

where

$$S = S^i \cap S^j \quad (\text{A.3})$$

Proof

Since there is a message transmission delay of one unit, it follows that S^i is in $G_i(\infty)$ and $G_j(\infty)$. Since $u_i(t)$ is $G_i(t+1)$ -measurable it follows $u_i(\infty)$ is $G_i(\infty)$ -measurable. Similarly $u_j(\infty)$ is $G_j(\infty)$ -measurable.

Consequently,

$$u_i(\infty) I(S'') = E[X | G_i(\infty) \cap G_j(\infty)] I(S'').$$

Similarly,

$$u_j(\infty) I(S'') = E[X | G_i(\infty) \cap G_j(\infty)] I(S'').$$

Hence,

$$u_i(\infty) I(S) = u_j(\infty) I(S)$$

■

To proceed further we need the following result:

Lemma A2.

Let $z_1, z_2, \dots, z_{q+1} = z_1$ be random vectors and $\underline{F}_1, \underline{F}_2, \dots, \underline{F}_q$ be σ -fields such that

$$z_i = E(z_{q+1} | \underline{F}_i) \quad i = 1, 2, \dots, q. \quad (A.4)$$

Then

$$z_1 = z_2 = z_3 = \dots = z_q \quad \text{a.s.} \quad (A.5)$$

Proof

We can assume that z_i are scalars, since by applying the same argument to each component we can generalize the result to random vectors. Suppose first that each z_i is square integrable. Since conditional expectation is the best mean square estimate and $z_i = E(z_{q+1} | \underline{F}_i)$, it follows that

$$E|z_{q+1}|^2 = E|z_i|^2 + E|z_{q+1} - z_i|^2 \quad i = 1, 2, \dots, q$$

Adding the above relations and using $z_{q+1} = z_1$ we get

$$0 = \sum_{i=1}^q E|z_{q+1} - z_i|^2.$$

Consequently $z_1 = z_2 = z_3 = \dots = z_q$. Thus, lemma A. holds for square integrable

random variable. To complete the proof of Lemma A2, for any number K let $z_i^K = \min \{z_i, K\}$. Then, by Jensen's inequality, $z_i = E\{z_{i+1} | \underline{F}_i\}$ implies

$$z_i^K \geq E\{z_{i+1}^K | \underline{F}_i\} \quad \forall i \quad (\text{A.6})$$

The last inequality implies

$$E z_1^K \geq E z_2^K \geq \dots \geq E z_q^K \geq E z_{q+1}^K = E z_1^K.$$

Consequently, (A.6) holds with equality.

Therefore, for $k_1 > k_2$

$$z_i^{k_1} - z_i^{k_2} = E(z_{i+1}^{k_1} - z_{i+1}^{k_2} | \underline{F}_i).$$

Since $z_i^{k_1} - z_i^{k_2}$ is bounded, it is square integrable, therefore

$$z_1^{k_1} - z_1^{k_2} = z_2^{k_1} - z_2^{k_2} = \dots = z_q^{k_1} - z_q^{k_2}.$$

Lemma 2 follows by letting $k_1 \rightarrow \infty$ and $k_2 \rightarrow \infty$.

Lemmae A1 and A2 can now be used to prove the following result

Lemma A3.

Suppose that

- (i) $i_1, i_2, \dots, i_{q+1} = i$, form a communication ring for S , and
- (ii) $1(S)$ is common knowledge for $G_1(\infty), G_2(\infty), \dots, G_q(\infty)$.

Then $u_i(\infty)$ agree on S , i.e.,

$$u_1(\infty)1(S) = u_2(\infty)1(S) = \dots = u_q(\infty)1(S) \quad \text{a.s.} \quad (\text{A.7})$$

Proof

By Lemma A1

$$\begin{aligned}
u_i(\infty) 1(S^{k+1}) &= E(X | G_i(\infty) \cap G_{i+1}(\infty)) = \\
&= E(u_{i+1}(\infty) | G_i(\infty) \cap G_{i+1}(\infty)) 1(S^{k+1}) = \\
&= E(u_{i+1}(\infty) 1(S^{k+1}) | G_i(\infty) \cap G_{i+1}(\infty))
\end{aligned} \tag{A.8}$$

By hypothesis (ii) $S \subset G_i(\infty)$ and $S \subset S^{k+1}$. Multiplication of both sides of (A.8) by $1(S)$ gives

$$u_i(\infty) 1(S) = E(u_{i+1}(\infty) 1(S) | G_i(\infty) \cap G_{i+1}(\infty)) \tag{A.9}$$

Eq.(A.9) and Lemma A2 imply

$$u_1(\infty) 1(S) = u_2(\infty) 1(S) \cdots = u_q(\infty) 1(S)$$

Lemma A3 can now be used to prove the following result.

Lemma A4

Under the hypothesis of Lemma A3

$$u_i(\infty) 1(S) = E(X | G_1(\infty) \cap G_2(\infty) \cdots \cap G_q(\infty)) 1(S) \tag{A.10}$$

Proof:

By (A.9)

$$u_i(\infty) 1(S) = E(X 1(S) | G_i(\infty) \cap G_{i+1}(\infty)) \tag{A.11}$$

By Lemma A3

$$u_i(\infty) 1(S) = u_j(\infty) 1(S),$$

thus, $u_i(\infty) 1(S)$ is common knowledge for $G_1(\infty), G_2(\infty), \dots, G_q(\infty)$. Taking conditional expectation with respect to $G_1(\infty) \cap G_2(\infty) \cdots \cap G_q(\infty)$ we obtain

$$\begin{aligned}
u_i(\infty) 1(S) &= E(X 1(S) | G_1(\infty) \cap G_2(\infty) \cap \cdots \cap G_q(\infty)) \\
&= E(X | G_1(\infty) \cap G_2(\infty) \cdots \cap G_q(\infty)) 1(S)
\end{aligned}$$

since by hypothesis $1(S)$ is common knowledge for $G_1(\infty), G_2(\infty), \dots, G_q(\infty)$. ■

The assertion of Theorem 1 now follows from Lemma A4 since $1(\Omega)$ is common knowledge for all persons. The estimate of each agent converges to $E(X | G_1(\infty) \cap G_2(\infty) \cdots \cap G_q(\infty))$.

APPENDIX B

Proof of Theorem 2

The information of team member i is described by the σ -field \underline{Y}_i . The σ -fields \underline{Y}_i evolve dynamically as follows:

$$\underline{Y}_i(t+1) = \underline{Y}_i(t) \vee \bigvee_{j \in [i]} \sigma(d(\underline{Y}_j(t))), \quad (i = 1, 2, \dots, \eta) \quad (\text{B.1})^*$$

with initial condition

$$\underline{Y}_i(0) = y_i(0) \quad (i = 1, 2, \dots, \eta) \quad (\text{B.2})$$

where $[i]$ is the set of team members with whom i communicates either directly or indirectly. By assumption all the team members belong to the same communication ring; thus, (B.1) can be written as

$$\underline{Y}_i(t+1) = \underline{Y}_i(t) \vee \bigvee_{j \neq i} \sigma(d(\underline{Y}_j(t))), \quad i = 1, 2, \dots, \eta \quad (\text{B.3})$$

Since $\underline{Y}_i(t) \uparrow \underline{Y}_i(\infty)$, it follows by the continuity of the decision rule d that

$$\lim_{t \rightarrow \infty} u_i(t) = u_i(\infty) \quad (\text{B.4})$$

Then equations (B.1) and (B.2) imply that for each k, j we have

$$\sigma(d(\underline{Y}_k(t))) \subset \underline{Y}_j$$

and

$$\sigma(d(\underline{Y}_k(\infty))) \subset \underline{Y}_j.$$

Then the agreement condition for rings implies that

$$u_1(\infty) = u_2(\infty) = u_3(\infty) = \dots = u_\eta(\infty) = d(\cap \underline{F}_i(\infty))$$

*Note that \vee is the join operation on σ -fields: $\underline{F}_1 \vee \underline{F}_2$ is the smallest σ -field containing \underline{F}_1 and \underline{F}_2 .

APPENDIX C

Proof of Proposition 1

At first we show that if

$$d(\underline{F}') = \max \{ d(\underline{G}) : \underline{G} \subset \underline{F}, \sigma(d(\underline{G})) \subset \underline{F}' \} \quad (\text{C.1})$$

then d satisfies the agreement condition for rings. Suppose that $\sigma(d(\underline{F}')) \subset \underline{G} \subset \underline{F}'$.

Because $\underline{G} \subset \underline{F}'$, $d(\underline{G}) \leq d(\underline{F}')$. Since $\sigma(d(\underline{F}')) \subset \underline{G}$, it is clear that $d(\underline{F}') \in \{d(\underline{H}) : \underline{H} \subset \underline{F}, \sigma(d(\underline{H})) \subset \underline{G}\}$. Thus, $d(\underline{F}') \leq d(\underline{G})$. The relation \leq is a partial order, consequently $d(\underline{G}) \leq d(\underline{F}')$ and $d(\underline{F}') \leq d(\underline{G})$ imply $d(\underline{F}') = d(\underline{G})$.

Hence d satisfies the agreement condition for pairs and in particular

$$d(\sigma(d(\underline{F}'))) = d(\underline{F}').$$

Suppose $\underline{F}_1 = \underline{F}_{\eta+1}$ and $\sigma(d(\underline{F}_K)) \subset \underline{F}_{K+1}$ for $1 \leq K \leq \eta$.

Then

$$d(\underline{F}_K) = d(\sigma(d(\underline{F}_K))) \leq d(\underline{F}_{K+1})$$

for each K , hence $d(\underline{F}_1) \leq d(\underline{F}_K) \leq d(\underline{F}_{\eta+1} = d(\underline{F}_1))$, and so $d(\underline{F}_1) = d(\underline{F}_K)$ for all K .

This shows that the decision rule d defined by Eq.(C.1) satisfies the agreement condition for rings.

Conversely, suppose that d satisfies the agreement condition for rings. Define the partial order \leq on $\{d(\underline{F}') : \underline{F}' \subset \underline{F}\}$ as follows: Write $d(\underline{F}_1) \leq d(\underline{F}_2)$ if and only if there is an integer $\eta \geq 1$ and σ -fields $\underline{G}_K \subset \underline{F}$, $1 \leq K \leq \eta$, such that $\sigma(d(\underline{F}_1)) \subset \underline{G}_1$, $\sigma(d(\underline{G}_K)) \subset \underline{G}_{K+1}$ and $d(\underline{G}_\eta) = d(\underline{F}_2)$. It is easy to see that $d(\underline{F}') \leq d(\underline{F}'')$ for all $\underline{F}' \subset \underline{F}$, (hence \leq is reflexive), and that $d(\underline{F}_1) \leq d(\underline{F}_2)$ and $d(\underline{F}_2) \leq d(\underline{F}_3)$ imply $d(\underline{F}_1) \leq d(\underline{F}_3)$ (hence \leq is transitive). Suppose $d(\underline{F}_1) \leq d(\underline{F}_2)$ and $d(\underline{F}_2) \leq d(\underline{F}_1)$. Then there are σ -fields $\underline{G}_K \subset \underline{F}$, $1 \leq K \leq \eta + m$, such that $\sigma(d(\underline{G}_{\eta+m})) \subset \underline{G}_1$, $\sigma(d(\underline{G}_K)) \subset \underline{G}_{K+1}$, $1 \leq K \leq \eta + m - 1$, $d(\underline{F}_1) = d(\underline{G}_{\eta+m})$, and $d(\underline{F}_2) = d(\underline{G}_\eta)$. The agreement condition implies that $d(\underline{G}_K) = d(\underline{G}_1)$ for all K , therefore $d(\underline{F}_1) = d(\underline{F}_2)$.

Consequently \leq is antisymmetric and so \leq is a partial order. Finally, if $\underline{G} \subset \underline{F}$ and $\sigma(d(\underline{G})) \subset \underline{F}'$, then $d(\underline{G}) \leq d(\underline{F}')$ by definition of \leq . Hence, $d(\underline{F}')$ is the maximum element of $\{d(\underline{G}) : \underline{G} \subset \underline{F}, \sigma(d(\underline{G})) \subset \underline{F}'\}$ with respect to \leq .

APPENDIX D

Proof of Proposition 2

Suppose $\delta_1, \delta_2 : \Omega \rightarrow U$ are \underline{F} -measurable. Define $\delta_1 \leq' \delta_2$ to mean either that $L(\delta_2) < L(\delta_1)$ or that $L(\delta_2) = L(\delta_1)$ and $\delta_1(\omega) \leq' \delta_2(\omega)$ for all ω . It is easy to see that \leq' so defined partially orders all \underline{F} -measurable decision functions. Suppose that $\underline{F}' \subset \underline{F}$ and δ is an \underline{F}' -measurable decision function. Since $d(\underline{F}') \in D(\underline{F}')$, by assumption $L(\delta) \leq L(d(\underline{F}'))$. If $L(\delta) = L(d(\underline{F}'))$, then $\delta \in D(\underline{F}')$ also, and $\delta(\omega) \leq' d(\underline{F}')(\omega)$ for all ω . It follows that $d(\underline{F}')$ maximizes $\{\delta : \sigma(\delta) \in \underline{F}\}$ with respect to \leq' . In particular, $d(\underline{F}')$ maximizes $\{d(\underline{G}) : \sigma(d(\underline{G})) \in \underline{F}', \underline{G} \subset \underline{F}\}$. Thus, Proposition 1 implies that d satisfies the agreement condition for rings.

APPENDIX E

Proof of Proposition 3

Assume that the condition

$$\sigma(d(\underline{F}_1)) \subset \underline{F}_2 \subset \underline{F}_1 \implies d(\underline{F}_1) = d(\underline{F}_2) \quad (\text{E.1})$$

is true. Then $\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2$ implies $\sigma(d(\underline{F}_1)) \subset \underline{F}_1 \cap \underline{F}_2 \subset \underline{F}_1$ which in turn implies $d(\underline{F}_1) = d(\underline{F}_1 \cap \underline{F}_2)$. Likewise $d(\underline{F}_2) = d(\underline{F}_1 \cap \underline{F}_2)$.

Hence, the condition

$$\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2 \implies d(\underline{F}_1) = d(\underline{F}_2)$$

is true.

Conversely, assume that

$$\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2 \implies d(\underline{F}_1) = d(\underline{F}_2) \quad (\text{E.2})$$

is true. Then $\underline{F}_2 \subset \underline{F}_1$ implies $\underline{F}_2 = \underline{F}_1 \cap \underline{F}_2$. Hence, $\sigma(d(\underline{F}_1)) \subset \underline{F}_2$ implies $\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_2 = \underline{F}_1 \cap \underline{F}_2$. Because of (E.2) it follows that $d(\underline{F}_1) = d(\underline{F}_2)$. Thus, the condition

$$\sigma(d(\underline{F}_1)) \subset \underline{F}_2 \subset \underline{F}_1 \implies d(\underline{F}_1) = d(\underline{F}_2)$$

is true.

APPENDIX F

Proof of Theorem 4

The proof of Theorem 4 proceeds in various steps: First we describe precisely the evolution of the estimation process according to each team member's model, and determine what each member's model predicts. Then, compare these predictions with what happens in reality. Both *Alpha* and *Beta* can describe the evolution of the estimation process according to their own view of the world as follows: Let a_t^i and β_t^i be the estimates of *Alpha* and *Beta* at time t according to i 's perception. ($i = \text{Alpha}, \text{Beta}$). Then,

$$a_t^i = d^i(A, \beta_1^i, \beta_2^i, \dots, \beta_{t-1}^i) \quad (\text{F.1})$$

$$\beta_t^i = d^i(B, a_1^i, a_2^i, \dots, a_t^i) \quad (\text{F.2})$$

where d^i denotes that the estimates are formed according to the rule d and the probability measure p^i induced by the distribution P^i on Ω . Equations (F.1) and (F.2) considered for all t and for all $\omega \in \Omega$ describe the evolution of the estimation process according to member i 's view of the world. To determine what *Alpha* and *Beta* predict about the outcome of the estimation process in terms of their own models consider the following σ -fields

$$\underline{F}_t^{iA} = \sigma(A, \beta_1^i, \beta_2^i, \dots, \beta_{t-1}^i) \quad (i = \text{Alpha}, \text{Beta}) \quad (\text{F.3})$$

$$\underline{F}_t^{iB} = \sigma(B, a_1^i, a_2^i, \dots, a_t^i)$$

The σ -fields \underline{F}_t^{iA} \underline{F}_t^{iB} describe the view of member i about the information available to *Alpha* and *Beta* after the initial observations have been taken and t tentative decisions have been exchanged. The σ -fields \underline{F}_t^{iA} and \underline{F}_t^{iB} evolve dynamically as follows:

$$\begin{aligned} \sigma(d^*(\underline{F}^B)) &\subset \underline{F}^{A^*} \cap \underline{F}^B \subset \underline{F}^A \\ \text{and } \sigma(d^*(\underline{F}^{A^*})) &\subset \underline{F}^{A^*} \cap \underline{F}^B \subset \underline{F}^A \end{aligned} \quad (\text{F.12})$$

Since d satisfies the agreement conditions, (F.12) implies that

$$d^*(\underline{F}^{A^*}) = d^*(\underline{F}^{A^*} \cap \underline{F}^B) = d^*(\underline{F}^B). \quad (\text{F.13})$$

Thus, both *Alpha* and *Beta* predict that the estimates will converge and agree after a finite number of steps.

In reality, the following is happening: At time $t = 1$ *Alpha's* estimate is

$$\bar{a}_1 = d^1(\bar{A}) \quad (\text{F.14})$$

(where \bar{A} is *Alpha's* observation). The message \bar{a}_1 is transmitted to *Beta*. *Beta* interprets this message according to his own view of the world, i.e., he considers that the realization \tilde{A} of A is such that

$$\bar{a}_1 = a_1^B = d^2(\tilde{A}). \quad (\text{F.15})$$

Furthermore, for a consistent interpretation of the data it is required that

$$P^B(\bar{a}_1 = a_1^B) > 0. \quad (\text{F.16})$$

At $t = 2$ *Beta's* estimate is

$$\bar{\beta}_1 = d^2(\bar{B}, \bar{a}_1) \quad (\text{F.17})$$

(where \bar{B} is *Beta's* observation), and this estimate is transmitted to *Alpha* who interprets it in terms of his own model, i.e., he considers that the realization \tilde{B} of B is such that

$$\bar{\beta}_1 = \beta_1^A = d^1(\tilde{B}, \bar{a}_1) \quad (\text{F.18})$$

For a consistent interpretation of the data it is required that

$$P^A(\beta_1^A = \bar{\beta}_1) > 0. \quad (\text{F.19})$$

In general, when *Alpha* receives message $\bar{\beta}_K$ he interprets it in terms of his own view of the world, i.e.,

$$\bar{a}_{K+1} = d^1(\bar{A}, \bar{\beta}_1, \bar{\beta}_2, \dots, \bar{\beta}_K) \quad (\text{F.20})$$

which he sends to *Beta*. For a consistent interpretation of all the messages received by

Alpha and Beta, it is required that at any time t

$$\begin{aligned} P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq t) &> 0 \\ \text{and } P^B(a_l^B = \bar{a}_l, 1 \leq l \leq t) &> 0 \end{aligned} \quad (\text{F.22})$$

The following result about the evolution of the probabilities of (F.22) is true.

Proposition F.1.

After a finite number of steps s^A , either

$$P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq s^A) = 0 \quad (\text{F.23})$$

or

$$P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq s^A) = 1 \quad (\text{F.24})$$

Moreover, for all $s > s^A$

$$P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq s) = P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq s^A). \quad (\text{F.25})$$

Similar results hold for $P^B(a_l^B = \bar{a}_l, 1 \leq l \leq t)$.

Proof

The result follows directly from the fact that convergence and agreement are predicted to occur in a finite number of steps by both models. The time s^A is given by (F.1).

Based on the proposition above we can complete the proof of Theorem 4 as follows:

If

$$\begin{aligned} P^A(\beta_l^A = \bar{\beta}_l, 1 \leq l \leq t) &> 0 \\ P^B(a_l^B = \bar{a}_l, 1 \leq l \leq t) &> 0 \end{aligned}$$

are true for all $t < s^A, s^B$ respectively, and (F.24) is true for both $P^A(\cdot)$ and $P^B(\cdot)$, then because of (F.13) and the rules by which the messages are interpreted

$$d^A(F_l^{AA}) = d^A(F_l^{AB}) = d^B(F_l^{BB}) = d^B(F_l^{BA})$$

for all $l \geq \max(s^A, s^B)$ and the estimates of Alpha and Beta agree asymptotically. If

on the other hand (F.23) is true at some time t for either *Alpha* or *Beta*, then, at that time *Alpha* or *Beta* realize that the sequence of received messages is impossible, or more reasonably, *Alpha* or *Beta* must conclude that the two models P^A and P^B are inconsistent.




Figure 1

APPENDIX D

FURTHER RESULTS ON THE CONSENSUS PROBLEM

FURTHER RESULTS ON THE CONSENSUS PROBLEM

Dr. David A. Castanon*

Prof. Demosthenis Teneketzis**

December, 1985

* ALPHATECH, Inc.
111 Middlesex Turnpike
Burlington, MA 01803

** Dept. of Electrical Engineering & Computer Science
University of Michigan
Ann Arbor, Michigan 48109

This work was sponsored by the Office of Naval Research under Contract
N00014-84-C-0845

ABSTRACT

In this paper, we develop additional results on the problem of reaching a consensus of opinion between two decisionmakers provided with different information. Specifically, we study the problem where the two decisionmakers may have different underlying probability models. We develop results characterizing the likelihood of a consensus being reached in terms of the nature of the inter-decisionmaker communications. We also study the problem when the decisionmakers are aware of the possibility that they may have different models. In this case, the decisionmakers can reach a deadlock state where neither decisionmaker can learn additional information from the consensus process, and they cannot reach a consensus decision. This surprising result indicates that incorporating human uncertainty in probability assessment into the consensus problem can lead to outcomes not anticipated in the general theory developed in refs. [1] - [7].

1. INTRODUCTION

The general problem of reaching a consensus of opinion among several decisionmakers provided with different information has received considerable attention in the recent literature [1]-[7]. The consensus problem consists of finding a decision, which, to each decisionmaker, is the correct decision according to a specific decision rule, given his information. Decisionmakers approach a consensus by exchanging tentative decisions among themselves, thus exchanging part of their information.

In [1]-[7], a Bayesian framework was developed for analyzing the consensus problem. Under the conditions that all decisionmakers share a common prior probability model, Aumann[1], Borkar and Varaiya [2], Tsitsiklis and Athans [3], Geanakoplos and Polemarchakis [4] and Washburn and Teneketzis [5] showed that decisionmakers would approach a consensus under mild regularity conditions on the communication pattern.

In subsequent papers [6][7], Teneketzis and Varaiya showed that relaxing the condition that all decisionmakers share a common probability model could lead to eventual disagreement. Specifically, they showed that, when the fact that each decisionmaker's probability model can be different is secret knowledge [12] (not available to any decisionmaker), the consensus process can reach a state of contradiction, thereby revealing that the underlying probability models were different.

In this paper, we examine in greater depth some of the issues raised by the results of Teneketzis and Varaiya [6],[7]. We limit our study to the case of two decisionmakers involved in the consensus problem. First, we study the question of how likely are the agreement or disagreement outcomes when the two decisionmakers have secret probability models. Then, we study the consensus problem where each decisionmaker can have multiple probability models, so that knowledge that the underlying probability models can be different is common knowledge. Based on this assumption, we develop a new Bayesian formulation of the consensus problem which is similar to the Bayesian formulation for games of incomplete information [8]. Using this formulation, we show that the general frameworks of Washburn and Teneketzis [5] and Teneketzis and Varaiya [7] can be extended to study issues of convergence and agreement in this problem.

The rest of this paper is organized as follows: In section 2, we describe the mathematical

framework which is used to study the consensus problem. In section 3, we develop additional results on the problem studied by Teneketzis and Varaiya [6], [7]. In section 4, we discuss the consensus problem with multiple models. Section 5 contains a discussion of the results.

2. PROBLEM FORMULATION

Throughout this paper, we will use the following stochastic decisionmaking model:

Let (Ω, \mathcal{F}) denote a measurable space, with \mathcal{F} denoting the σ -field of measurable events*. Let $\{P^i, i \in I\}$ denote a family of probability measures on this measure space. The set I is assumed to be a discrete set, finite or countable, with the discrete topology. The measure space (Ω, \mathcal{F}) represents the uncertainty present in the decision problem.

There are two decisionmakers (DM) in the consensus problem. DM 1 (2) has a personal probability model $(\Omega, \mathcal{F}, P^{i_1})$ ($(\Omega, \mathcal{F}, P^{i_2})$), where i_1 and i_2 are selected from the index set I . In addition, each DM has a probability distribution over I , representing his beliefs that the other DM is using a particular model, as follows:

Denote by $I \times I$ the event space of all possible combinations of models for the two DMs, with the product σ -field $2^I \times 2^I$. Let P_1, P_2 denote probability measures on this space. DM 1's (2's) initial private information concerning the pair of probability measures $\{P^{i_1}, P^{i_2}\}$ used by the DM's is represented by the σ -field H_1 (H_2), where H_1 is generated by atoms h^1_j of the form

$$h^1_j = \{i_1\} \times K_j, \text{ where } I \supset K_j \quad (1)$$

Similarly, H_2 is generated by atoms h^2_i of the form

$$h^2_i = K_i \times \{i_2\}, \text{ where } I \supset K_i \quad (2)$$

Eqs. 1 and 2 mean that each DM is provided with private information concerning his own probability model, plus the information that the other agent's model belongs to a specific set of models.

In addition to private information concerning the personal probability models used by each DM, each DM receives private information concerning the true event which occurs in the measurable

* The reader unfamiliar with the concepts of measure theory should refer to [10].
space (Ω, F) . This information is represented by finite-valued F -measurable functions

$$y^i: \Omega \longrightarrow Y^i \quad (3)$$

Let Y^i in F denote the σ -field induced by y^i . Then, Y^i is a finite σ -field.

We can combine the measurable spaces (Ω, F) and $(I \times I, 2^I \times 2^I)$ to form the product space $(\Omega \times I \times I, F \times 2^I \times 2^I)$. On this space, define the measures Π_1, Π_2 as follows: Let $F \in F$. Then,

$$\Pi_1(F, i_1, i_2) = P_1(i_1, i_2) P^1_1(F) \quad (4)$$

$$\Pi_2(F, i_1, i_2) = P_2(i_1, i_2) P^2_2(F) \quad (5)$$

Π_1, Π_2 are probability measures on $(\Omega \times I \times I, F \times 2^I \times 2^I)$ because P^i are probability measures for each i .

For two σ -fields A, B , define $A \vee B$ to be the smallest σ -field containing both A and B . Similarly, define $A \wedge B$ to be the largest σ -field contained in both A and B . Define Q_1, Q_2 to be the restrictions of Π_1, Π_2 to the σ -field

$$\underline{E} = \{Y^1 \vee Y^2\} \times \{H_1 \vee H_2\}. \quad (6)$$

The decision rule used by each decisionmaker is a map from σ -fields A into decision functions, which depends on the probabilities Q_i as in [5]. We assume that the decision rules are \underline{E} -measurable. For the problems considered in this paper, the decision rules will be of the form

$$d_i(F) \in \arg \min_{E \in Q_i} \{J(\omega, u) | F\} \quad (7)$$

$$u \in U$$

$$d_2(F) \in \arg \min_{u \in U} E^{Q_2} \{ J(\omega, u) | F \} \quad (8)$$

for $F \in \mathcal{L}$ where U is the space of allowable decisions, and the subscript i is used to denote the expectation according to the probability distribution Q_i . Note that multiple solutions to eqs. 7 and 8 can occur. Usually, there will be a tie-breaking procedure for selecting $d_1(F)$, $d_2(F)$. We assume that U is compact, and that

$$J: \Omega \times U \longrightarrow [0, \infty) \quad (9)$$

is a continuous function of u for each $\omega \in \Omega$. As in [5], let $\sigma(d_i(A))$ be the σ -field generated by the decision rule d under probability Q_i when the available information σ -field is A .

The consensus process can now be described. Each DM receives initially one measurement y_i . Based on this measurement and his probability model, each DM computes a tentative decision according to a decision rule d_i and communicates it to the other DM. Then, each DM sequentially interprets the other DM's decision, revises his own decision due to the acquired knowledge, and communicates his new decision to the other DM. This process creates a sequence of information σ -fields $F_1(n), F_2(n)$ evolving in the lattice of sub- σ -fields of \mathcal{E} . With this process, a sequence of decisions is generated.

Let $u_1(n)$ ($u_2(n)$) denote the value of the n th communication of DM 1 (DM 2), selected as a function of the information available to him according to his decision rule. For (ω, i_1, i_2) in $\Omega \times I \times I$, we say that the DMs reach a consensus (agreement) at (ω, i_1, i_2) if and only if

$$\lim_{n \rightarrow \infty} u_1(n) = \lim_{n \rightarrow \infty} u_2(n) \quad (10)$$

The above framework includes the formulations of [1] -[7] as special cases. In [1]-[5], the common probability model formulation can be captured in the above framework by letting the set I be the singleton set $\{1\}$. In this case, the probabilities P_1, P_2 are trivial, and the remaining probabilistic framework corresponds to the general framework presented in [5]. The formulation of [6] can be

captured by selecting P_1, P_2 to be purely diagonal measures, of the form

$$P_1(i_1, i_2) = 0 = P_2(i_1, i_2) \text{ if } i_1 \neq i_2. \quad (11)$$

and the selection of i_1, i_2 is such that $i_1 \neq i_2$. In this case, each decisionmaker is convinced that the other decisionmaker will use the same probability model as he does. However, the initial models selected for each player may be different.

3. GENERICITY AND CONTINUITY OF CONSENSUS: THE SECRET MODEL PROBLEM

In this section, we analyze the model of [6] to determine how likely are agreement or disagreement outcomes. We separate our results into two cases: the case when the decision variables are continuous, and the case when the decision variables are discrete. In order to specialize the formulation of section 2 to the problems investigated in [6] and [7], we make the following assumptions:

A1. There exists some $A \in \{Y^1 \vee Y^2\}$ such that

$$P^1(A) \neq P^2(A), \quad (12)$$

A2. The beliefs of each DM concerning the other DM's probability model, P_1 and P_2 , satisfy eq. 11.

Assumption A1 guarantees that the differences in the DM's models are detectable with the available observation. Assumption A2 specifies that the knowledge that the models may be different is secret knowledge to each DM.

3.1 Continuous decision variables

When the decision space U is a continuous space, we make the following additional assumption:

A3. U is a convex subset of R^n , and $J(\omega, u)$ is a strictly convex, differentiable function of u for each ω .

Assumption A3 guarantees that there exist unique solutions to eqs. 7 and 8. Note that A1 - A3 are satisfied by the model in [6], since the decisionmakers exchange the conditional probability of an event (X) occurring given their information. In this case,

$$J(\omega, u) = (I\{\omega \in X\} - u)^2, \quad (13)$$

so it satisfies A3. The assumptions in [6] concerning the different probability models of the decisionmakers correspond to A1 and A2.

In order to characterize the likelihood of agreement or disagreement results, we need the following definitions: Let Π be the space of all probability distributions on $(\Omega, \{Y^1 \vee Y^2\})$. Since $\{Y^1 \vee Y^2\}$ is a finite σ -field, Π is a simplex in R^n , where n is the cardinality of the atoms of $\{Y^1 \vee Y^2\}$. Alternatively, Π can be viewed as a subset of R^{n-1} with positive Lebesgue measure.

Definition: A result is said to be generic in Π if and only if the set $\{\Pi \in \Pi \mid \text{result is not true for } \Pi\}$ has zero $n-1$ -dimensional Lebesgue measure.

Proposition 1: Under assumptions A1-A3, if for all $B \supset A$, either

- i. $u_1(\omega) = \arg \min_{u \in U} E_{\pi_1}\{J(\omega, u) \mid B\} \in U^0$, or
- ii. $u_2(\omega) = \arg \min_{u \in U} E_{\pi_2}\{J(\omega, u) \mid B\} \in U^0$

where U^0 is the interior of U , it is generic that, for some $\omega \in \Omega$, a contradiction will be reached in the process of consensus.

Proof: Without loss of generality, assume condition i holds. Consider any instance when a tentative decision is sent from DM 1 to DM 2. Denote that decision as u_1 , and the information σ -field available to player 1 as F_1 . Then, u_1 is an F_1 -measurable random variable satisfying

$$u_1(\omega) = \arg \min_{u \in U} \text{Epi}_1 \{ J(\omega, u) \mid F_1 \}. \quad (14)$$

In order for a contradiction not to occur, DM 2 must be able to interpret $u_1(\omega)$ in terms of his own probability model; that is,

$$u_{21}(\omega) = \arg \min_{u \in U} \text{Epi}_2 \{ J(\omega, u) \mid F_{21} \}, \quad (15)$$

where F_{21} is DM 2's perception of the information available to DM 1. Note that F_{21} is a coarser σ -field than $\{Y^1 \vee Y^2\}$, hence it is also a finite field. Let U_{21} denote the following subset of U :

$$U_{21} = \{ u \in U \mid u = \arg \min_{u \in U} \text{Epi}_2 \{ J(\omega, u) \mid B \} \text{ for some } B \in F_{21}, \text{ for some } \sigma\text{-field } F_{21} \text{ satisfying } \{Y^1 \vee Y^2\} \supseteq F_{21} \supseteq Y^1 \}$$

Note that U_{21} is a finite set, since each σ -field F_{21} is finite and there are only a finite number of σ -fields satisfying the inclusion conditions. A necessary condition for the consensus process not to reach a contradiction is $u_1(\omega) \in U_{21}$ for all $\omega \in \Omega$.

Let f_i denote the atoms of $\{Y^1 \vee Y^2\}$. Let $F_1(\omega)$ denote the atom of F_1 containing ω , and $f(\omega)$ denote the atom of $\{Y^1 \vee Y^2\}$ containing ω . Select $\omega \in A$ such that $P^1_1(f(\omega)) \neq P^1_2(f(\omega))$. Such a ω exists by assumption A2. Then,

$$\begin{aligned} u_1(\omega) &= \arg \min_{u \in U} \{ \text{Epi}_1 \{ \text{Epi}_1(J(\omega, u) \mid \{Y^1 \vee Y^2\}) \mid F_1 \} \\ &= \arg \min_{u \in U} \text{Epi}_1 \left\{ \frac{\int_{v \in f(\omega)} J(v, u) P^1_1(dv)}{\int_{v \in f(\omega)} P^1_1(dv)} \mid F_1(\omega) \right\} \end{aligned}$$

Define $J(f_i, u)$ as

$$\underline{J}(f_i, u) = \frac{\int_{v \in f_i} J(v, u) P^{i1}(dv)}{P^{i1}(f_i)} \quad (16)$$

Then,

$$u_1(\omega) = \operatorname{argmin}_{u \in U} \left\{ \frac{\sum_{f_i \subset F_1(\omega)} P^{i1}(f_i) \underline{J}(f_i, u)}{P^{i1}(F_1(\omega))} \right\} \quad (17)$$

By assumption, the minimizing value is in the interior of U ; a necessary and sufficient condition characterizing $u_1(\omega)$ is

$$\sum_{F_1(\omega) \supset f_i} P^{i1}(f_i) \frac{\partial}{\partial u} I(f_i, u_1(\omega)) = 0. \quad (18)$$

Since $u_1(\omega)$ must belong to U_{21} , this means, for some $\alpha \in U_{21}$, $\Delta \in F_{21}$,

$$\sum_{\Delta \supset f_i} P^{i2}(f_i) \frac{\partial}{\partial u} I(f_i, \alpha) = 0. \quad (19)$$

For each α, Δ the set of P^{i2} in Π satisfying eq. 19 has $n-1$ dimensional Lebesgue measure 0, since eq. 19 imposes a linear constraint on P^{i2} . Since there is a finite number of α in U_{21} and Δ in F_{21} , the set of P^{i2} in Π satisfying eq. 19 for some α, Δ also has $n-1$ dimensional Lebesgue measure 0, which implies that a contradiction is generic for some ω in A . q.e.d.

The results of proposition 1 can be understood in terms of the example below:

Example 1: Let $\Omega = [0,2] \times [0,3]$, and let F denote the Borel sets in Ω . Let the σ -field Y^1 be defined by the atoms $\{\omega | \omega \in a_1\}$, $\{\omega | \omega \in a_2\}$, and the σ -field Y^2 be defined by the atoms $\{\omega | \omega \in b_1\}$,

$\{\omega|\omega \in b_2\}$ and $\{\omega|\omega \in b_3\}$, where a_i, b_j are defined in figure 1. Define probability models P^1, P^2 as:

$P^1(A) = \mu(A)/6$, where μ is two-dimensional Lebesgue measure

$P^2(A) = \mu(A \cap b_1)/36 + \mu(A \cap b_2)/24 + 7\mu(A \cap b_3)/72$.

P^1 is uniformly distributed over Ω , while P^2 is uniformly distributed conditioned on b_i , but has

$P^2(b_1) = 1/6; P^2(b_2) = 1/4; P^2(b_3) = 7/12$.

Let $i_1 = 1$, while $i_2 = 2$, so that DM 1 uses probability model P^1 , while DM 2 uses probability model P^2 . The decision rule used by the DM's is defined by eq. 13, where the event is event X in figure 1.

As noted in [6], when $\omega \in a_1 \cap b_3$, and DM 1 communicates first, there is eventual agreement, although the DMs have very different reasons for reaching that agreement. In order to show that disagreement is generic, we will show that arbitrarily small perturbations to P^2 will result in disagreement. Specifically, let

$P(b_1) = (2 + \epsilon_1)/12; P(b_2) = (3 + \epsilon_2)/12; P(b_3) = (7 - \epsilon_1 - \epsilon_2)/12$.

As in [6], DM1's first communication is $u_1 = .5$. In order for this value not to be a contradiction, either

$$3\epsilon_1 + 2\epsilon_2 = 0 \quad (20a)$$

or

$$\epsilon_1 + 2\epsilon_2 = 4. \quad (20b)$$

The two-dimensional Lebesgue measure of the set of all ϵ_1, ϵ_2 satisfying eqs. 20 a or b is 0, since it is the union of two lines. Hence, for almost all choices of ϵ_1, ϵ_2 , a contradiction will be reached in the first communication.

The reason for the genericity of disagreement in proposition 1 is that, although each DM can observe only a finite number of observation values, he can communicate a continuous number of decisions. This enables the other DM to detect differences in the probability models. Conditions i or ii in proposition 1 guarantee that the announced decisions will vary with small differences in probability models.

3.2 Discrete Decision Variables

In this subsection, we assume that the space U is discrete. Let f_i denote an atom of $\{Y^1 \vee Y^2\}$. We define a metric on Π as follows:

For $P^1, P^2 \in \Pi$,

$$d(P^1, P^2) = \max_{f_i} |P^1(f_i) - P^2(f_i)|. \quad (21)$$

This metric is equivalent to the Euclidean metric on Π .

Definition: An agreement or disagreement result is said to be continuous in Π at P^1, P^2 if agreement or disagreement continues to hold for all P^1, P^2 in a neighborhood of P^1, P^2 .

Assume in addition:

A4. For any $A \in \{Y^1 \vee Y^2\}$, there exists unique u_1, u_2 in U such that

$$u_1 = \arg \min_{u \in U} E_{P^1} \{J(\omega, u) \mid \omega \in A\}$$

$$u_2 = \arg \min_{u \in U} E_{P^2} \{J(\omega, u) \mid \omega \in A\}$$

With this assumption, we have the following characterization of agreement or disagreement outcomes:

Proposition 2. Under assumptions A1, A2 and A4, if agreement occurs for P^1, P^2 , it is continuous in Π . If disagreement occurs for P^1, P^2 , it is continuous in Π .

Proof: Without loss of generality, assume that P^1, P^2 result in agreement. Denote by $F_1(n)$ ($F_2(n)$) DM 1's (DM 2's) sequence of σ -fields generated in the consensus process. Each one of these fields is coarser than $\{Y^1 \vee Y^2\}$, hence finite. For any time interval n , the atoms of $F_1(n)$ and $F_2(n)$ are

elements of $\{Y^1 \vee Y^2\}$. For any $A \in \{Y^1 \vee Y^2\}$, define the function

$$L(P^1, A) = \left\{ \frac{\sum_{f_1 \in A} P^1(f_1) \underline{J}(f_1, u)}{P^1(A)} \right\} \quad (22)$$

This is a continuous function in Π . Because of assumption A4, we can find a $\delta_1(A)$ such that, for $d(P^1, P^1) \leq \delta_1(A)$,

$$\begin{aligned} \operatorname{argmin}_{u \in U} \left\{ \frac{\sum_{f_1 \in A} P^1(f_1) \underline{J}(f_1, u)}{P^1(A)} \right\} &= \\ \operatorname{argmin}_{u \in U} \left\{ \frac{\sum_{f_1 \in A} P^1(f_1) \underline{J}(f_1, u)}{P^1(A)} \right\} & \end{aligned}$$

A similar result can be established in terms of $\delta_2(A)$ for DM 2's decisions. Select σ as

$$\delta = \min_{\{Y^1 \vee Y^2\} \ni A} \{ \delta_1(A), \delta_2(A) \}. \quad (23)$$

This minimum exists because there are only a finite number of A in $\{Y^1 \vee Y^2\}$. This choice of δ guarantees that the exchanged sequence of decisions and the σ -fields inferred by the other DM are the same for all probabilities P^1, P^2 satisfying

$$\begin{aligned} d(P^1, P^1) &\leq \delta \\ d(P^2, P^2) &\leq \delta, \end{aligned}$$

thereby completing the proof. q.e.d.

The result of proposition 2 depends critically on assumption A4. However, an argument similar to the proof of proposition 1 establishes the following result.

Proposition 3. If, for each atom $f_i \in \{Y^1 \vee Y^2\}$,

$$I(f_i, u) \neq I(f_i, v) \text{ if } u \neq v,$$

then Assumption A4 is generic in Π .

Proof: If assumption A4 does not hold for P^1 , there must exist a set $A \in \{Y^1 \vee Y^2\}$ and $u, v \in U$, $u \neq v$, such that

$$E_{P^1} \{J(\omega, u) \mid \omega \in A\} = E_{P^1} \{J(\omega, v) \mid \omega \in A\}.$$

This implies that

$$\sum_{f_i \in A} P^1(f_i) \underline{J}(f_i, u) = \sum_{f_i \in A} P^1(f_i) \underline{J}(f_i, v).$$

Since $I(f_i, v) \neq I(f_i, u)$ for any atom of $\{Y^1 \vee Y^2\}$, this implies that the set of P^1 which satisfy this equation has Lebesgue measure 0 in Π . A similar argument for P^2 completes the proof. q.e.d.

Example 2: Let (Ω, \mathcal{F}) , Y_1 , Y_2 , and X be defined as in example 1. Define probability models P^3 , P^4 as:

$$P^3(A) = 2\mu(A \cap a_1)/15 + \mu(A \cap a_2)/5, \text{ where } \mu \text{ is two-dimensional Lebesgue measure,}$$

$$P^4(A) = \mu(A \cap a_2)/5 + 4\mu(A \cap a_1 \cap b_1 \cap X)/95 + 12\mu(A \cap a_1 \cap b_1 \cap \bar{X}) + 12\mu(A \cap b_2 \cap a_1)/95 + 36\mu(A \cap a_1 \cap b_3 \cap X)/95 + 12\mu(A \cap a_1 \cap b_3 \cap \bar{X})/95$$

where \bar{X} is the complement of X in Ω . Rather than work with the unconditional probabilities, what is important is to evaluate the conditional probabilities of events given available information. Thus, P^3 is uniformly distributed conditioned on a_i , with

$$P^3(a_1) = .4; P^3(a_2) = .6.$$

P^4 has the same distribution as P^3 on a_2 , but differs on a_1 , as

$$P^4(X|a_1 \cap b_1) = .25; P^3(X|a_1 \cap b_1) = .5 \quad (24a)$$

$$P^4(X|a_1 \cap b_2) = P^3(X|a_1 \cap b_2) = .75 \quad (24b)$$

$$P^4(X|a_1 \cap b_3) = .50; P^3(X|a_1 \cap b_3) = .25 \quad (24c)$$

$$P^4(X|a_1 \cap (b_2 \cup b_3)) = .6; P^3(X|a_1 \cap (b_2 \cup b_3)) = .5 \quad (24d)$$

$$P^4(X|a_1) = 10/19; P^3(X|a_1) = .5 \quad (24e)$$

Let $U = \{0,1\}$. Let

$$J(\omega, 0) = .53 \text{ if } \omega \in X,$$

$$= 0 \text{ if } \omega \in \bar{X},$$

$$J(\omega, 1) = 0 \text{ if } \omega \in X$$

$$= .47 \text{ if } \omega \in \bar{X}.$$

With this definition of J , the optimal decision for DM 1 given an information set A is given by

$$\begin{aligned} u_1 &= 1 \text{ if } P^1_1(\omega \in X | A) > .47 \\ &= 0 \text{ otherwise.} \end{aligned} \quad (25)$$

The same decision rule is optimal for DM 2, using the probability P^1_2 .

Assume $i_1 = 3$, while $i_2 = 4$, that DM 1 exchanges his decision first, and that the DM's alternate in exchanging decisions. As in example 1, assume that $\omega \in a_1 \cap b_3$. From eqs. 24e and 25, DM 1's initial decision is 1. That is, $u_1(1) = 1$. If DM 1's information had been a_2 , his decision would have been $u_1 = 0$. Hence, according to DM 1, he has signalled $\omega \in a_1$ to DM 2.

According to DM 2's probability model, $u_1(1) = 1$ implies $\omega \in a_1$. Hence, DM 2 believes $\omega \in a_1 \cap b_3$. His optimal decision is $u_2(1) = 1$, because of eq. 24c. According to DM 2, his decision has signalled $\omega \in b_2 \cup b_3$, because of eqs. 24a and 24b.

Because of eqs. 24 a,b,c, DM 1 interprets $u_2(1) = 1$ to mean $\omega \in a_1 \cap (b_1 \cup b_2)$. His optimal decision is $u_1(2) = 1$. Hence, a consensus has been reached at $u_1(2) = u_2(2) = 1$. However, they have reached this agreement for the wrong reasons, since DM 1 believes $\omega \in a_1 \cap (b_1 \cup b_2)$, whereas in actuality, $\omega \in a_1 \cap b_3$! Note that any changes in either P^3 or P^4 which would change the numbers in eqs. 24 a-c by less than .02 would continue to result in agreement.

Suppose that the order of communication is reversed, so that DM 2 communicates first. The optimal decision $u_2(1) = 1$. Note that $P^4(X | \omega \in b_2) < .45$, $P^4(X | \omega \in b_1) < .2$. Hence, DM 2 believes he has signaled $\omega \in b_3$.

According to DM 1, he interprets $u_2(1) = 1$ to mean $\omega \in b_3$. Hence, he believes $\omega \in a_1 \cap b_3$. His optimal decision, according to eq. 24c, is $u_1(1) = 0$. This decision cannot be understood by DM 2, because he expected $u_1(1) = 1$ whether DM 1 knew a_1 or a_2 . Hence, the DMs have reached a contradiction. Note that this contradiction will be reached even if P^i_1 or P^i_2 are modified by .02. Hence, the disagreement outcome is also continuous.

The above results illustrate that, when the decision spaces are discrete, small discrepancies in the probability models of the decisionmakers will not affect the consensus process. They also show that the set of pairs of probability models for which consensus occurs has positive Lebesgue measure in Π , unlike the result in the continuous decision case of the previous section. However, the set of pairs of probability models for which contradictions occur also has positive Lebesgue measure. Hence, contradictions are common phenomena in the consensus process.

The question still remains: How does the consensus process proceed once a contradiction is encountered? Such a contradiction reveals that the basic assumption that $P_1(i_1, i_2) = P_2(i_1, i_2) = 0$ if $i_1 \neq i_2$, is violated. In the next section, we present a plausible model for this process, and study its implications.

4. CONSENSUS PROBLEMS WITH MULTIPLE PROBABILITY MODELS

When the decisionmakers in the consensus process have different subjective views of the world, and these differences are secret knowledge, the results of [6]-[7] show that a contradiction outcome is often reached whereby the existence of these differences becomes common knowledge. At this point, our model of how the consensus process proceeds is that each decisionmaker models statistically the types of probability models which the other decisionmaker may be employing, and acquires information through the consensus process concerning the possible models used by the other DM, and the uncertainty in the event space Ω . Within this framework, we investigate convergence and agreement issues for two cases: When the statistics of the types of probability models are common knowledge, and when these statistics are secret knowledge. The analysis is based on the mathematical formulation developed in section 2, where the probability distributions P_1, P_2 represent the statistics used by each DM.

We make the following assumptions:

A5: P_1, P_2 are common knowledge.

A6: The decision rule d satisfies the agreement condition

Recall that P_1, P_2 are the subjective probabilities on the space of possible model pairs (i_1, i_2) .

Assumption A5 implies that the subjective statistical distribution of possible probability models for each decisionmaker is known to the other decisionmaker, and this fact is common information. Note that we do not assume that these distributions are equal. This allows DM 1 to believe he has a different range of possible decision models than DM 2 has, and viceversa.

The decision rule d is said to satisfy the agreement condition if, whenever $G_1 \supseteq G_2 \supseteq \sigma(d(G_1))$, then $d(G_1) = d(G_2)$. The agreement condition implies that, if a decision is based on information which is common to the information σ -fields G_2 and G_1 , then knowledge of either G_2 or G_1 would result in the same decision. In [5] and [7], a sufficient condition was developed to characterize when a decision rule satisfies the agreement condition. When specialized to our model, this condition can be stated as:

Proposition 4: Assume that there is a total order $<$ on U . Let $D_1(\omega), D_2(\omega)$ denote the set of solutions of eqs. 7, 8 for each ω . If $u_1(\omega), u_2(\omega)$ are selected to be the minimal elements in $D_1(\omega), D_2(\omega)$ respectively, then the decision rules d_1, d_2 satisfy the agreement condition.

Proof: See [5], [7].

Under assumptions A5 and A6, we can prove the following result:

Proposition 5: Under Assumptions A5 and A6, if $i_1 = i_2$ and $\omega \in \Omega$ is such that the consensus process reveals that $i_1 = i_2$ to both decisionmakers, then the decisionmakers reach a consensus for ω .

Proof: Without loss of generality, let $i_1 = i_2 = 1$. The consensus process starts with the initial information σ -fields

$$G_1(0) = Y^1 \times H_1 \quad (26a)$$

$$G_2(0) = Y^2 \times H_2 \quad (26b)$$

After each communication is heard, each decisionmaker learns additional information. The evolution of information of each decisionmaker can be described by the evolution of a dynamical system in the lattice of sub- σ -fields of \mathcal{E} , as in [5]:

$$G_1(n+1) = G_1(n) \vee \sigma(d_2(G_2(n))) \quad (27 a)$$

$$G_2(n+1) = G_2(n) \vee \sigma(d_1(G_1(n))) \quad (27 b)$$

$$(\text{or } G_1(n+1) = G_1(n) \vee \sigma(d_2(G_2(n+1))) \quad (27 c)$$

$$G_2(n+1) = G_2(n) \vee \sigma(d_1(G_1(n))), \quad (27 d)$$

depending on whether communications are simultaneous or staggered),

where $d_i(G)$ is the decision rule of eqs. 7 or 8 applied to the atoms of the σ -field G . The lattice operations are \vee and \wedge , where $A \vee B$ represents the coarsest σ -field containing both A and B , and

$A \wedge B$ is the finest σ -field contained in both A and B . Because of A5, the evolutions indicated in eq. 27 are common knowledge. Note that these dynamical systems are evolving on a lattice of finite fields, and that they generate a strictly increasing sequence of σ -fields. Hence, after some finite time t , a limit must be reached such that, for all $s > t$,

$$G_1(s) = G_1 = G_1 \vee \sigma(d_2(G_2)) \quad (28 a)$$

$$G_2(s) = G_2 = G_2 \vee \sigma(d_1(G_1)). \quad (28 b)$$

Eqs. 28 establish that the consensus process converges to a limit; that is,

$$\lim_{n \rightarrow \infty} u_1(n) = u_1^*; \quad \lim_{n \rightarrow \infty} u_2(n) = u_2^*. \quad (29)$$

In addition, eq. 28 implies that

$$G_1 \supseteq \sigma(d_2(G_2)) \quad (30 a)$$

$$G_2 \supseteq \sigma(d_1(G_1)) \quad (30 b)$$

Furthermore, the fields generated by a decision rule are contained in the information available for decisions. That is,

$$G_2 \supseteq \sigma(d_2(G_2)) \quad (31 a)$$

$$G_1 \supseteq \sigma(d_1(G_1)) \quad (31 b)$$

Hence,

$$G_2 \supseteq G_2 \wedge G_1 \supseteq \sigma(d_2(G_2)) \quad (32 a)$$

$$G_1 \supseteq G_2 \wedge G_1 \supseteq \sigma(d_1(G_1)) \quad (32 b)$$

By assumption, the decision rules satisfy the agreement condition. Hence, by eq. 32,

$$\begin{aligned} d_1(G_1) &= d_1(G_1 \wedge G_2) \\ d_2(G_2) &= d_2(G_1 \wedge G_2) \end{aligned} \quad (33)$$

Select $\omega \in \Omega$ such that $i_1 = 1 = i_2$ is common knowledge. Since the σ -fields are increasing, there is an atom g containing ω in $G_1 \wedge G_2$ of the form $g = (A, 1, 1)$, where $\Omega \supseteq A$. Since g is an element of both limiting fields G_1 and G_2 , it follows that

$$\begin{aligned} d_1(g) &= \arg \min_{u \in U} E^{Q_1} \{ J(\omega, u) | g \} \\ &= \arg \min_{u \in U} \{ P_1(1, 1) E_{P_1} \{ J(\omega, u) | \omega \in A \} \}. \end{aligned} \quad (34)$$

Similarly, since $i_2 = 1$ in g ,

$$\begin{aligned} d_2(g) &= \arg \min_{u \in U} E^{Q_2} \{ J(\omega, u) | g \} \\ &= \arg \min_{u \in U} \{ P_2(1, 1) E_{P_2} \{ J(\omega, u) | \omega \in A \} \}. \end{aligned} \quad (35)$$

Note that the functions being minimized are simple multiples of each other. Thus, $d_2(g) = d_1(g)$ for all such g . Coupled with eq. 33, this completes the proof. q.e.d.

According to the above proposition, even if the decisionmakers have the same probability model, consensus is not guaranteed unless it becomes common knowledge that $i_1 = i_2$. On the other hand, the condition that $\omega \in \Omega$ is such that the consensus process reveals that $i_1 = i_2$ is sufficient but not necessary for reaching a consensus. These points are illustrated by the following example.

Example 3: Let (Ω, F) , Y_1, Y_2, X, P^3, P^4, U and $J(\omega, u)$ be defined as in example 2. Consider the decision rule defined in eqs. 7 and 8. Assume that $i_1 = i_2 = 3$, so that both decisionmakers have the same probability model. Assume further that there exists a probability distribution Q on $\{3, 4\}$ such that

$$P_1(i_1=j, i_2=k) = P_1(i_1=j, i_2=k) = Q(j)Q(k). \quad (36)$$

That is, each decisionmaker believes that the other DM's decision model is selected independently from a known statistical population, where the set of possible models was $\{P^3, P^4\}$. Furthermore, the statistics of the selection are known identically to both decisionmakers. However, the precise model selected is private knowledge provided to each decisionmaker.

Assume $Q(3) = .1$. As in example 2, we assume that the DMs alternate exchanging tentative decisions. As in example 2, DM 1's first decision is $u_1(1) = 1$, and this signals that $\omega \in a_1$ to DM 2. This decision does not reveal whether $i_1 = 3$ or $i_1 = 4$, because if $i_1 = 3$ or 4, the same decision would be made.

At this point in the consensus process, DM 2 knows that $\omega \in a_1 \cap b_3$ whether $i_1 = 3$ or 4. Hence, by eq. 8, since $i_2 = 3$, the optimal decision is $u_2(1) = 0$. In order to identify the information signalled by this decision, we must examine the optimal decisions corresponding to the possible information sets that DM 2 could have, from DM 1's perspective. These decisions are:

If $i_2 = 3, \omega \in a_1 \cap b_1$, then $u_2(1) = 1$ by eq. 24 a

If $i_2 = 3, \omega \in a_1 \cap b_2$, then $u_2(1) = 1$ by eq. 24 b

If $i_2 = 3, \omega \in a_1 \cap b_3$, then $u_2(1) = 0$ by eq. 24 c

If $i_2 = 4, \omega \in a_1 \cap b_1$, then $u_2(1) = 0$ by eq. 24 a

If $i_2 = 4, \omega \in a_1 \cap b_2$, then $u_2(1) = 1$ by eq. 24 b

If $i_2 = 4, \omega \in a_1 \cap b_3$, then $u_2(1) = 1$ by eq. 24 c.

Hence, DM 1 knows that either $i_2 = 3, \omega \in a_1 \cap b_3$ or $i_2 = 4, \omega \in a_1 \cap b_1$. According to eq. 7, his optimal decision is selected as

$$\begin{aligned}
u_1(2) &= \arg \min_{u \in U} E^{Q_1} \{ J(\omega, u) | F \} \\
&= \arg \min_{u \in \{0,1\}} \{ Q(3) [(.25)(.53)I\{u=0\} + (.75)(.47)I\{u=1\}] + \\
&\quad Q(4) [(.5)(.53)I\{u=0\} + (.5)(.47)I\{u=1\}] \} \quad (37)
\end{aligned}$$

where $I\{\}$ is the indicator function, and a constant scaling factor has been omitted. It is easy to see that, for $Q(3) < .12$, $u_1(2) = 1$. Note that, if $i_1 = 4$, eq. 37 becomes

$$\begin{aligned}
u_1(2) &= \arg \min_{u \in \{0,1\}} \{ Q(4) [(2/3)(.53)I\{u=0\} + (2)(.47)I\{u=1\}] + \\
&\quad Q(3) [(3)(.53)I\{u=0\} + (3)(.47)I\{u=1\}] \} \quad (38)
\end{aligned}$$

So, $u_1(2)$ should be 0 if $i_1 = 4$. Therefore, $u_1(2) = 1$ signals that $i_1 = 3$ to DM 2.

DM 2 now knows $i_1 = i_2 = 3$, and $\omega \in a_1 \cap b_3$. As before, his optimal decision is $u_2(2) = 0$. This decision does not convey any additional information to DM 1, because the decision $u_2(2) = 0$ did not depend on the information $i_1 = 3$. Since DM 1 obtains no additional information, his tentative decision continues to be $u_1(2) = 1$, and the two decisionmakers agree that an agreement cannot be reached. The common information which forms the basis for this disagreement can be summarized in the atom

$$\{(\omega, i_1, i_2) | (\omega \in a_1 \cap b_3, i_1 = i_2 = 3), \text{ or } (\omega \in a_1 \cap b_1, i_1 = 3, i_2 = 4)\}.$$

Consider now the same problem, but assume that DM 2 communicates first. Then, since $\omega \in b_3$ and $P^3(X | b_3) = .55$, then $u_2(1) = 1$. DM 1 observes a_1 and receives $u_2(1) = 1$; by the same argument as above, he concludes that DM 2 has observed b_3 and uses either model 3 or 4. His decision is $u_1(1) = 0$.

DM 2 is aware that DM 1 knows both a_1 and b_3 , so when he receives $u_1(1) = 0$, he interprets this to mean that $i_1 = 3$ and $\omega \in a_1$ and b_3 . Hence, he communicates $u_2(2) = 0$. This reveals that $i_2 = 3$, and the decisionmakers reach a consensus.

The results of example 3 are rather surprising. Unlike the cases studied in [6] or [7], there is no unmodeled secret information present in this consensus process. Indeed, both DM's are actually using the same probability model; furthermore, they have identical probability distributions over the class of probability models, and this is common knowledge! Nevertheless, a disagreement outcome occurs. This implies that even admitting the possibility that the other DM can have a different subjective probability model than your own is sufficient to prevent reaching a consensus. The reason for this effect is the difference in the probability distributions used by each DM in eqs. 7 and 8 when one DM is unable to identify the probability model used by the other DM.

How likely is it that the conditions of proposition 5 are met? Our analysis of the previous section can be extended to establish the following propositions:

Proposition 6. Suppose that the decision space U was continuous, and that assumption A3 was true. Assume $i_1 = i_2$. Assume additionally that, for all $B \in \mathcal{E}$

$$\text{i. } u_1(\omega) = \arg \min_{u \in U} E^{Q_1}\{J(\omega, u) | B\} \in U^0$$

$$\text{ii. } u_2(\omega) = \arg \min_{u \in U} E^{Q_2}\{J(\omega, u) | B\} \in U^0$$

where U^0 is the interior of U . Then, the outcome that the consensus process will reveal that $i_1 = i_2$ is generic in Π^k , where k is the cardinality of I .

Proposition 7: If the decision space U is discrete, $i_1 = i_2$, and for any $A \in \mathcal{E}$, there exists unique u_1, u_2 in U such that

$$u_1 = \arg \min_{u \in U} E^{Q_1}\{J(\omega, u) | A\}$$

$$u_2 = \arg \min_{u \in U} E^{Q_2}\{J(\omega, u) | A\}$$

then the outcome that a consensus process reveals that $i_1 = i_2$ for a specific ω is continuous in Π^k .

The proof of these propositions follows directly the proof of propositions 1 and 2, and will not be reproduced here. Essentially, proposition 6 is based on the fact that the set of probability models for which a continuous decision fails to discriminate among a finite set of models has zero Lebesgue measure in the space of all possible probability models. Under the assumptions of proposition 7, one can show that the sequence of σ -fields generated in the consensus process does not change with small perturbations in the set of individual probability models.

When $i_1 \neq i_2$, it is possible to show, by arguments similar to those leading to eqs. 27-32 that the sequence of decisions $\{d_1(G_1(n))\}$ and $\{d_2(G_2(n))\}$ will converge to $d_1(G_1)$ and $d_2(G_2)$ respectively. However, since the probabilistic models of DM 1 and DM 2 are not the same, whether or not a consensus is reached depends on the event $\omega \in \Omega$ and the order of communication. If g is an atom of $G_1 \wedge G_2$ containing $\omega \in \Omega$, then a consensus will be reached if $d_1(g) = d_2(g)$.

The above results were based on the assumption that the underlying statistical models P_1 and P_2 used by each DM are common knowledge. When these models differ, and this fact is secret knowledge, and the decision processes of DM 1 and DM 2 are consistent with their own beliefs, then the consensus process reaches one of three different outcomes after a finite number of communications:

1. A consensus is reached,
2. DM 1 and DM 2 realize that their underlying statistical models are inconsistent,
3. DM 1 and DM 2 agree to disagree because they cannot gather any additional information from the consensus process.

In order to establish this, we must describe the evolution of the decision processes according to each DM's subjective decision model, and determine what each DM's model predicts. Then, we compare the predicted communications with the actual communications heard. Let $u_{11}(n)$, $u_{12}(n)$ denote the decisions of DM 1 and DM 2 at stage n according to DM 1's subjective decision model. Similarly, let $u_{21}(n)$, $u_{22}(n)$ denote the decisions of DM 1 and DM 2 at stage n according to DM 2's subjective decision model. Then, according to DM 1's view,

$$u_{11}(n) = d_{11}(y^1, u_{12}(1), \dots, u_{12}(n-1)) \quad (39 a)$$

$$u_{12}(n) = d_{12}(y^2, u_{11}(1), \dots, u_{11}(n-1)), \quad (39 \text{ b})$$

where d_{1j} denotes the decisions formed by the decision rule d according to the probability measure $\Upsilon \Pi_1$. Similarly, according to DM 2's view,

$$u_{21}(n) = d_{21}(y^1, u_{12}(1), \dots, u_{12}(n-1)) \quad (40 \text{ a})$$

$$u_{22}(n) = d_{22}(y^2, u_{11}(1), \dots, u_{11}(n-1)), \quad (40 \text{ b})$$

where d_{2j} denotes the decisions formed by the decision rule d according to the probability measure $\Upsilon \Pi_2$. Equations 39 and 40 describe the evolution of the consensus process according to DM 1's and DM 2's perception, respectively.

To determine what DM 1 and DM 2 predict about the outcome of the decision processes in terms of their own perceptions, we define four sequences of information σ -fields, representing DM 1's actual knowledge ($G_{11}(n)$), DM 1's belief of DM 2's knowledge ($G_{12}(n)$), DM 2's actual knowledge ($G_{22}(n)$) and DM 2's belief of DM 1's knowledge ($G_{21}(n)$). These fields evolve under communications as:

$$G_{11}(n+1) = G_{11}(n) \vee \sigma(d_{12}(G_{12}(n))) \quad (41 \text{ a})$$

$$G_{12}(n+1) = G_{12}(n) \vee \sigma(d_{11}(G_{11}(n))) \quad (41 \text{ b})$$

$$G_{21}(n+1) = G_{21}(n) \vee \sigma(d_{22}(G_{22}(n))) \quad (41 \text{ c})$$

$$G_{22}(n+1) = G_{22}(n) \vee \sigma(d_{21}(G_{21}(n))). \quad (41 \text{ d})$$

with initial conditions

$$G_{11}(0) = G_{21}(0) = Y^1 \times H_1 \quad (41 \text{ e})$$

$$G_{22}(0) = G_{12}(0) = Y^2 \times H_2. \quad (41 \text{ f})$$

As before, these evolutions occur in a lattice of σ -fields where the maximal element is a finite σ -field. Hence, repeating the logic of the proof of proposition 4 establishes that the consensus process will reach steady-state after a finite number of iterations. Denote this finite number as T .

To establish the type of outcomes possible, we must examine the consensus process closely. At stage 1, DM 1's decision is $u_{11}(1) = d_{11}(y_1)$. This message is transmitted to DM 2, who must interpret this message according to his own subjective decision model. That is, he must find realizations (i', y^1) of possible models and observation values such that $u_{11}(1) = d_{21}(y^1)$. For a consistent interpretation, one must have

$$\text{Prob}\{ (y^1 | u_{11}(1) = d_{21}(y^1), i_1 = i') > 0. \quad (42)$$

If this is not possible, DM 2 will discover that the decision models are inconsistent, leading to outcome 2. Otherwise, DM 2 selects $u_{22}(1) = d_{22}(y^2, u_{11}(1))$.

At this stage, DM 1 must interpret consistently the communications heard from DM 2. As before, he must find realizations (j', y^2) of possible models and observation values such that $u_{22}(1) = d_{12}(y^2, u_{11}(1))$. For a consistent interpretation, one must have

$$\text{Prob}\{ (y^2 | u_{22}(1) = d_{12}(y^2, u_{11}(1)), i_2 = j') > 0. \quad (43)$$

Define $P^1(n, i)$, $P^2(n, j)$ as follows:

$$P^1(n, i) = P^i\{ (y^1 | u_{11}(j) = d_{21}(y^1, u_{22}(1), \dots, u_{22}(j)), i_1 = i, \text{ for all } j < n \} \quad (44 a)$$

$$P^2(n, j) = P^j\{ (y^2 | u_{22}(j) = d_{12}(y^2, u_{11}(1), \dots, u_{11}(j)), i_2 = j, \text{ for all } j < n \} \quad (44 b)$$

It is easy to see that, for each $i \in I$, $P^1(n, i)$ and $P^2(n, i)$ are monotone decreasing sequences in n . Since the consensus process reaches steady-state after a finite number of communications (for $n \geq T$), there are three possible outcomes:

1. There exist no i or j such that both $P^1(T, i)P^1(i, i_2) > 0$ and $P^2(T, j)P^2(j, j_1) > 0$ (45)

2. There exist one i and one j such that

(46)

$$P^1(T,i)P_1(i,i_2) > 0$$

$$P^2(T,j)P_2(i_1,j) > 0$$

3. There exist more than one i or j such that

(47)

$$P^1(T,i)P_1(i,i_2) > 0$$

$$P^2(T,j)P_2(i_1,j) > 0$$

When eq. 45 holds, the inconsistencies among the statistical models P_1 and P_2 are detected in the consensus process. It is possible that the true probability model used by DM 1 was not considered possible in DM 2's subjective distribution. In this case, for some n , either $d_{12}(T) \neq d_{22}(T)$ or $d_{21}(T) \neq d_{11}(T)$.

When equation 46 holds, either a consensus outcome will be reached for $\omega \in \Omega$, or an inconsistency in the underlying probability models P^i and P^j will be discovered. If $i = j$, this is the case analyzed in [1]-[5]; in this case, the results of [5] guarantee that $d_{11}(T) = d_{12}(T) = d_{21}(T) = d_{22}(T)$, so a consensus outcome is reached. If $i \neq j$, this is the case analyzed in [6] and [7]. In this case, two types of outcomes are possible: either $d_{12}(T) \neq d_{22}(T)$ or $d_{21}(T) \neq d_{11}(T)$, so that an inconsistency in models P^i and P^j is detected, or both $d_{11}(T) = d_{21}(T)$ and $d_{22}(T) = d_{12}(T)$. For the second outcome, the results in [6] and [7] imply also that $d_{11}(T) = d_{22}(T)$.

When eq. 47 holds, there is residual ambiguity in both the statistical models and the underlying probability models. As discussed in example 3, it is possible to have $d_{12}(T) = d_{22}(T)$, $d_{21}(T) = d_{11}(T)$, and $d_{11}(T) \neq d_{22}(T)$ for all n , in addition to the other two outcomes. In this case, no inconsistencies have been discovered between either the statistical models P_1 and P_2 , or the underlying probability models P^i_1 and P^i_2 . Rather, the decisionmakers have reached a stage where no additional information will be exchanged in the consensus process. At this stage, the DMs agree that a consensus cannot be reached, and discontinue the process.

The above discussion has established the following proposition:

Proposition 8: Assume A6 is satisfied, and that the differences between P_1 and P_2 are secret knowledge to both decisionmakers. Then, the consensus process will reach one of three possible

outcomes after a finite number of communications:

1. A consensus is reached,
2. DM 1 and DM 2 realize that their underlying statistical models are inconsistent,
3. DM 1 and DM 2 agree to disagree because they cannot gather any additional information from the consensus process.

A result similar to Proposition 8 was obtained in [7]. However, when each DM considers a set of possible underlying probability models for the other DM, as is the case in this section, the consensus process can result in an outcome not predicted by the model of [7]; namely, the DMs can agree to disagree even though their underlying probability models have not been established as inconsistent with each other. This point is illustrated by the following example.

Example 4: Let (Ω, F) , Y_1 , Y_2 , X , P^1 , P^3 , P^4 be defined as in examples 2 and 3. Let $J(\omega, u)$ be defined as in eq. 13, and let $U = [0, 1]$. Let $I = \{1, 3, 4\}$ be the set of possible probability model indices. Consider the decision rule defined in eqs. 7 and 8. Assume that $i_1 = 1$, and $P_1(1, 1) = .3$, $P_1(1, 3) = 0$, $P_1(1, 4) = .3$, so that DM 1 believes DM 2 is using either model 1 or model 4 with equal probability. Assume $i_2 = 4$, and $P_2(1, 4) = 0$, $P_2(3, 4) = .1$, $P_2(4, 4) = .1$, so that DM 2 assumes that DM 1 is using either model 3 or model 4 with equal probability. Assume that $\omega \in a_1 \cap b_3$ and DM 1 communicates first.

The first tentative decision of DM 1 is $u_{11}(1) = .5$. According to DM 2, if DM 1 was using model P^3 , then $u_{21}(1) = .5$ when $\omega \in a_1$, and .25 when $\omega \in a_2$. If DM 1 was using model P^4 , then $u_{21}(1) = .25$ for any ω . Hence, DM 2 believes $i_1 = 3$, and $\omega \in a_1 \cap b_3$. According to his own model, P^4 , DM 2's communication is $u_{22}(1) = .5$.

In communicating $u_{11}(1) = .5$, DM 1 believes that he has signaled that $i_1 = 1$ and $\omega \in a_1$. Hence, DM 1 expects $u_{12}(1) = .5$ if $i_2 = 1$ and $\omega \in b_1$, or if $i_2 = 4$ and $\omega \in b_3$. Consequently, DM 1 chooses

$$\begin{aligned} u_{11}(2) &= \frac{P^1(X|a_1 \cap b_1) P^1(a_1 \cap b_1) P_1(1, 1) + P^1(X|a_1 \cap b_3) P^1(a_1 \cap b_3) P_1(1, 4)}{P^1(a_1 \cap b_1) P_1(1, 1) + P^1(a_1 \cap b_3) P_1(1, 4)} \\ &= .375 \end{aligned}$$

Since models P^3 and P^1 have the same distribution conditioned on a_1 , it is also true that $u_{21}(2) =$

.375. Since DM 2 believes he already knows $i_1 = 3$, and $\omega \in a_1 \cap b_3$, he learns no additional information, so his decision continues to be $u_{22}(2) = .5$. This decision conveys no additional information to DM 1, so the consensus process stalls at this point, and both decisionmakers agree that a consensus cannot be reached.

5. CONCLUSION

In this paper, we have studied the problem of reaching a consensus in a group of decisionmakers by exchanging tentative decisions using a Bayesian framework. When the decisionmakers have different probability models and the existence of those differences is secret knowledge, the results of Teneketzis and Varaiya [6], [7] characterized all possible outcomes of the consensus process into two types of outcomes:

1. Reaching a consensus decision for the group,
2. Reaching a contradiction.

The results of section 2 shed additional insight concerning how likely each of these outcomes is. By defining the concepts of a generic outcome and a continuous outcome, we have shown that, when the decision space is continuous-valued and some regularity conditions are met, reaching a contradiction is a generic outcome. In contrast, when the decision space is discrete-valued, both outcomes are continuous, so that small deviations in probability models result in the same outcomes.

One of the limitations of the results of Teneketzis and Varaiya is their assumption that knowledge that probability models could be different is secret knowledge to the decisionmakers, although in fact the probability models are different. If the decisionmakers are humans, subject to various biases and inaccuracies in evaluating probabilities [9], knowledge that there can exist differences in probability models is best modeled as common knowledge. In section 4, we developed a Bayesian framework whereby this knowledge is represented as common knowledge, and the specific individual probability models are represented as private information for each decisionmaker. In this framework, the consensus process serves both to reveal information concerning the probability model of each decisionmaker, as well as information concerning the problem uncertainty. A surprising result is that, even when the probability models of the decisionmakers are identical, and selected

independently from identical probability distributions, there are two possible outcomes:

1. A consensus was reached.
2. A point was reached where both decisionmakers, on the basis of common information, agree that a consensus cannot be reached.

The second outcome has not been predicted by the previous formulations [1] - [7]. Indeed, it seems to contradict the title of Geanakoplos and Polemarchakis [4], "We can't disagree forever." Our results in section 4 show that merely admitting the possibility that the probability models are different is sufficient to generate the second outcome. Again, we characterize how likely these outcomes are for both continuous-valued and discrete-valued decision spaces U .

In conclusion, we have shown that, in our Bayesian framework, when the decisionmakers bring human biases and inaccuracies in probability assessments into the consensus process, a consensus may not be reached even if the decisionmakers share the same probability model. The results depend explicitly on the Bayesian formulation for incorporating uncertainty concerning the other decisionmaker's true probability model. A different formulation, similar to Kreps and Wilson's formulation for sequential games [11], could be developed whereby each decisionmaker considers only the most likely interpretation of the results as the basis for selecting his tentative decisions. The merits of each formulation rest ultimately in their ability to help us understand the behavior of humans in consensus decisionmaking.




Fig. 1 a. Information Fields for each Decisionmaker




Fig. 1 b. Event X in the example

REFERENCES

1. Aumann, R.J., "Agreeing to disagree," Annals of Statistics, Vol. 4 (6), 1976, pp. 1236-1239.
2. Borkar, V. and P. Varaiya, "Asymptotic agreement in distributed estimation," IEEE Trans. on Automatic Control, Vol. AC-27, (3), 1982, pp. 650-655.
3. Tsitsiklis, J.N. and M. Athans, "Convergence and asymptotic agreement in distributed decision problems," IEEE Trans. Auto. Control, Vol. AC-29 (1), 1984, pp. 42-50.
4. Geanakoplos, J.D. and H.M. Polemarchakis, "We can't disagree forever," Journal of Economic Theory, Vol. 26, 1982, pp. 192-200.
5. Washburn, R.B. and D. Teneketzis, "Asymptotic agreement among communicating decisionmakers," Stochastics, vol. 13, 1984, pp. 103-129.
6. Teneketzis, D. and P. Varaiya, "Consensus in distributed estimation with inconsistent beliefs," Systems and Control Letters, 4, 1984, pp. 217-221.
7. Teneketzis, D. and P. Varaiya, "Consensus in distributed estimation," to appear in Advances in Signal Processing, V. Poor, ed., 1986 (also appendix C in this report).
8. Harsanyi, J., "Games with incomplete information played by Bayesian players," Management Science, Vol. 14, 1967, pp. 159-182.
9. DeGroot, M.H., "Reaching a Consensus," Journal of the American Statistical Association, Vol. 69, no. 345, March 1974.
10. Halmos, P., Measure Theory, Springer Verlag, New York, 1970.
11. Kreps, D.M. and R. Wilson, "Sequential Equilibria," Technical Rep. 340, IMSSS, Stanford University, December 1980.

12. Levine, P. and J. P. Ponsard, "The Values of Information in some Nonzero Sum Games,"
International Journal of Game Theory, Vol. 6, No. 4

APPENDIX E

INFORMATIONAL ASPECTS OF A CLASS OF SUBJECTIVE
GAMES OF INCOMPLETE INFORMATION

E-0

141

**INFORMATIONAL ASPECTS OF A CLASS OF SUBJECTIVE
GAMES OF INCOMPLETE INFORMATION**

Dr. Demosthenis Teneketzis*

Dr. David A. Castanon**

October, 1985

*** Dept. of Electrical Engineering & Computer Science
University of Michigan
Ann Arbor, Michigan 48109**

****ALPHATECH, Inc.
111 Middlesex Turnpike
Burlington, MA 01803**

ABSTRACT

Subjective games of incomplete information are formulated where some of the key assumptions of Bayesian games of incomplete information are relaxed. The issues arising because of the new formulation are studied in the context of a class of non-zero-sum two-person games, where each player has a different model of the game. Two types of games are investigated: static games and infinitely-repeated games. It is shown that in the static game counterintuitive outcomes of the game occur because of the different beliefs of the players, and that these outcomes indicate to the players that their models were different. When the game is repeated infinitely often, it is shown that this repetition can alleviate the differences in the models of the players and lead to eventual cooperation. When multiple equilibrium solutions are present, the effect of various bargaining theories on the outcome of the game is investigated. It is shown that, depending on the bargaining model, the players may agree on the outcome of the game or they may realize that they have different models.

SECTION 1: INTRODUCTION

Game theory is the mathematical science which studies decisionmaking in situations of potential conflict among decisionmakers. The requirements of formal game theory are strict regarding the ruler of the game and the portrayal of exogenous uncertainty. Due to these requirements, there are many strategic situations which cannot be initially modeled as games because players lack information about available strategies, utility functions or outcomes resulting from various strategies.

Specifically, the key requirements of formal game theory are:

- A1. The rules of the game are common information [15] to all players of the game
- A2. Exogenous uncertainty is portrayed by objective probabilities which are common knowledge to all players.
- A3. Players are fully committed to a priori strategies
- A4. Players are rational.

As Game theory developed, attempts were made to relax some of these assumptions. Assumption A3 was a consequence of the normalization principle of Von Neumann [12]; Aumann and Maschler [1] were the first to point out via a simple counterexample the inappropriateness of the normalization principle under certain conditions; since then, considerable developments followed by relaxing the assumption of prior commitment [2]-[6].

Harsanyi[7] and Aumann-Maschler et al. [8] pointed out that in some military problems, players may lack full information about the payoff functions of other players, or about the physical facilities and strategies of other players, or even about the amount of information that other players have about the various aspects of the game situation. Thus, Harsanyi [7] first relaxed assumption A1 and formulated and developed models of games of incomplete information. Harsanyi modeled the incomplete information as an exogenous random move (Nature's move) to select among possible games; he also assumed that the outcomes of this move have a (subjective) probability distribution which is common knowledge to all players. Considerable progress has been achieved in the theory of games of incomplete information using Harsanyi's original formulation (see [8]-[10] and references therein.)

A restriction in Harsanyi's formulation is the requirement of common knowledge of the probability distribution of Nature's move. In many strategic situations (especially in noncooperative games), this distribution is subjectively assessed by each player, and subject to individual biases and inaccuracies [27]. In this paper, we formulate a class of games, which we call Subjective Games of Incomplete Information, which relaxes Harsanyi's requirements of common knowledge. Specifically, we allow each person to have his own subjective probability distribution of nature's move; in addition, each person believes his subjective distribution is common information, whereas it is actually secret information [15]. As a consequence, requirements A1 and A2 are relaxed, and requirement A4 is modified in the sense that each player is considered to be rational within his/her own subjective view of the game.

Various interesting issues arise because of our formulation:

- Q1. How are equilibrium strategies defined for subjective games?
- Q2. How do these equilibrium strategies relate to the equilibrium strategies of the games studied so far?
- Q3. Does repetition of the game result in cooperation as in the case of the games studied so far (e.g. [11])? Does repetition of the game alleviate differences in the subjective assessments of the players and allow players to agree on an equilibrium strategy?
- Q4. Is it possible to characterize the set of all equilibria for repeated subjective games?

To understand some of these questions, we shall consider a special class of games, namely 2×2 two-person non-zero sum games of incomplete information where the payoff matrices have a special structure.

The rest of this paper is organized as follows: In section 2, we present the model for subjective games, and briefly discuss games of incomplete information and point out the differences between Harsanyi's model and our model. In section 3, we study static subjective non-cooperative games of incomplete information. In section 4, we consider infinitely repeated non-cooperative subjective games of incomplete information. Conclusions are presented in section 5.

SECTION 2: FORMULATION OF SUBJECTIVE GAMES OF INCOMPLETE INFORMATION

We shall develop our theory of subjective games based on the following key assumptions:

S1. Players have different probability assessments on nature's move.

S2. Each player thinks that the other players' assessments are the same as his.

S3. Players are Bayesian.

S4. Each player is rational within his own subjective view of the game.

Assumption S2 implies that the rules of the game are not common knowledge to all the players, since each player thinks that the other players' assessments are the same as his, yet this may not be true. Assumptions S1 and S2 were previously used in the context of distributed estimation and detection [14].

More precisely, let b_i represent the private information of player i about the game. This information relates to the outcome of nature's move. In dealing with incomplete information, each player takes a Bayesian approach. That is, each player assigns a subjective probability distribution P_i

$$P_i = P_i(b_1, b_2, \dots, b_i, \dots, b_n)$$

to nature's move and attempts to maximize the mathematical expectation of his own payoff J_i in terms of this probability distribution. Furthermore, each player i assumes that $P_i = P_j$ for all j , whereas in the actual game, P_i and P_j may be different.

Comparing the mathematical model described above with Harsanyi's formulation, we note that Harsanyi also assumes that each player assigns a subjective probability distribution P_i to nature's move; although P_i and P_j may differ, all the distributions P_i are assumed to be common knowledge to all players. In our formulation, any difference in subjective probabilities is secret information. Moreover, each player is unaware that he has secret information.

SECTION 3:

STATIC SUBJECTIVE NON-COOPERATIVE GAMES OF INCOMPLETE INFORMATION

3.1 Problem formulation

We consider the following static two-person non-zero sum game. Nature selects one of two games with the following payoff matrices:

Game 1:

	σ	τ	
λ	(a,a)	(c,b)	(3-1)
μ	(b,c)	(d,d)	

Game 2:

	σ	τ	
λ	(b,b)	(d,a)	(3-2)
μ	(a,d)	(c,c)	

We further assume that

$$a > c > b > d \quad (3-3)$$

$$b+c > a+d \quad (3-4)$$

Player 1 can choose action λ or μ and player 2 can choose action σ or τ . Note that, because of

(3-3), each player has a dominant strategy in each one of the two games. So far, the statement of the problem and the assumptions (3-3)-(3-4) are essentially the same as in [15]. However, contrary to [15], we now assume that the two players have a different probability assessment of nature's move. Let r be the true probability that nature selects game 1. Let p, q be player 1's and player 2's assessments of this event respectively. Assume $p > 1/2, q < 1/2$.

We will consider this problem under four different types of information that a player may receive:

1. No information: In this case, none of the players is informed about the outcome of nature's move.
2. Public information: In this case, both players are informed about the outcome of nature's move.
3. Private information: In this case, one player is informed about the outcome of nature's move, whereas the other player is not. Moreover, this distribution of information is common knowledge.
4. Secret information: In this case, one player is informed about the outcome of nature's move whereas the other player is uninformed. Moreover, the uninformed player is unaware that his opponent is informed, and the informed player knows this.

The rational strategies in each of these situations are:

1. No Information

In this case, player 1 plays λ and player 2 plays τ . The payoffs of the two players are:

$$J_1^0 = rc + (1-r)d, \quad (3-5)$$

and

$$J_2^0 = rb + (1-r)a. \quad (3-6)$$

2. Public information

In this case, player 1 plays λ in Game 1 and μ in game 2. Player 2 plays σ in game 1 and τ in game 2. Thus, the payoff of the players is

$$JB_1 = JB_2 = ra + (1-r)c. \quad (3-7)$$

Define the value of information as follows: V_i , the value of information to player i , is the payoff of player i when he knows the outcome of nature's move minus the payoff of player i when no player is informed about the outcome of nature's move.

In this case, the value of public information for players 1 and 2 is given by

$$VB_1 = r(a-c) + (1-r)(c-d) \quad (3-8)$$

$$VB_2 = (2r-1)a + (1-r)c - rb \quad (3-9)$$

3. Private Information

3a. Assume at first that player 1 is the informed player. Then he plays σ in game 1 and μ in game 2. Player 2 plays τ . The payoffs of the two players are

$$JP_1 = c \quad (3-10)$$

$$JP_2 = rb + (1-r)c \quad (3-11)$$

3b. If player 2 is the informed player, then he plays σ in game 1 and τ in game 2. If

$$p < (a-d)/(a+c-b-d) = p^* \quad (3-12)$$

then player 1 will play μ . Otherwise, he will play λ . The expected payoffs for player 1 are then

$$JP_1 = rb + (1-r)c \quad (3-13)$$

$$JP_2 = ra + (1-r)d \quad (3-14)$$

respectively.

The payoff for player 2 is

$$J^P_2 = c \text{ (corresponding to } \mu) \quad (3-15)$$

$$J^P_2 = a \text{ (corresponding to } \lambda) \quad (3-16)$$

The value of information for the two players is:

$$V^P_1 = (1-r)(c-d) \quad (3-17)$$

$$V^P_2 = c - rb - (1-r)a \text{ if } p < p^* \\ = r(a-b) \text{ otherwise} \quad (3-18)$$

4. Secret Information

Assume at first that player 1 is secretly informed about the outcome of the chance move.

Then, he plays λ in game 1 and μ in game 2. Player 2 plays τ . The payoffs of the two players are

$$J^C_1 = c \quad (3-19)$$

$$J^C_2 = rb + (1-r)c \quad (3-20)$$

The value of secret information to player 1 in this case is

$$V^S_1 = (1-r)(c-d) \quad (3-21)$$

Assume now that player 2 is secretly informed. Then, he plays σ in game 1 and τ in game 2.

Player 1 plays λ . The payoffs of the two players are

$$J^C_1 = r a + (1-r) d \quad (3-22)$$

$$J^C_2 = a \quad (3-23)$$

The value of secret information to player 2 in this case is

$$vS_2 = r(a-b) \quad (3-24)$$

Let us discuss some interesting features of the solutions of these games. At first, note that each payoff bimatrix is symmetric, hence in each one of the two games, the players are interchangeable. Thus, one expects that for the classical Bayesian game, in the case of public or secret information, the behavior of the informed and the uninformed player will be independent of who is the informed and who is the uninformed player. For example, in the case of private or secret information, if player 1 were the uninformed player and played λ , we would expect that if the situation were reversed and player 2 became the uninformed player he would play σ . Also, in the case where no player was informed about the outcome of the chance move, the dominant strategies would be (λ, σ) or (μ, τ) . Consequently, the value of private, secret or public information would be the same for both players. It can be easily checked that this is indeed the case when $p=q=r$. However, this behavior is not observed when each player has his own subjective model of the game. When player 1 is privately informed about the chance move, player 2 always chooses τ (the second column); on the other hand, if player 2 is privately informed about the outcome of the chance move, player 1 does not always play μ (the second row). When player 1 is the secretly informed player, player 2 always plays τ (second column); if player 2 is the secretly informed player, player 1 always plays λ (first row). When no player is informed about the outcome of the chance move, the outcome of the game is (λ, τ) . These facts indicate that the value of private and secret information is now different for each player, as is evident from the analysis above.

For the class of games considered in this section, the value of public, private and secret information differs from player to player, whereas in the classical Bayesian framework, this value does not depend on who is the informed and who is the uninformed player. This phenomenon is due to the differences in the initial probability assessments of the incomplete information.

Another interesting observation follows from the previous results. Consider the case where player 2 is privately informed, $p < p^*$, and $r = 1/2$. Then, the value of information for player 2 is given by

$$v^P_2 = c - .5b - .5a$$

if $c < .5a + .5b$, the value of private information for player 2 is negative! On the other hand, the gain for player 1, the uninformed player, is equal to $.5(b-d)$ which is positive. Thus, for the class of symmetric games considered in this paper, we have a case where the value of private information is negative for the informed player and the uninformed player benefits from the situation! This phenomenon never occurs for this class of games in the classical Bayesian framework, where if the value of private information is negative for the informed player, the uninformed player cannot benefit either [15]. Even more surprising in this case is the fact that the informed player wants to use his private information, whereas the uninformed player wishes that the informed player acted as if he were not informed!

The reason for all these counterintuitive results and the differences between the subjective game results and the classical Bayesian game results is that each player evaluates the game as well as the behavior of his opponent in the game in terms of his own model and acts accordingly. Such subjective evaluations lead to behavior which would never occur in the classical Bayesian formulation as evidenced by the previous analysis.

One issue that naturally arises in these games is the following: How do the players involved in the game interpret its outcome? Do they realize that they have different models? If neither player is informed about the outcome of the chance move player 1 expects that player 2 will use strategy σ and player 2 expects that player 1 will use strategy μ . At the end of the game, each player finds out that the outcome is the opposite of what he expected. Since each player assumes that his opponent is rational, both players conclude that they have different models. Similar phenomena occur if one of the players is either secretly or privately informed.

In the case of secret information, the secretly-informed player discovers at the end of the game that his opponent's perception of the game is different from his. On the other hand, the uninformed player may never discover that his opponent has a different perception of the game, or he may not be able to interpret his opponent's move in terms of his own model in which case he can conclude that either his opponent has a different model of the game, or his opponent has secret information.

In the case of private information, the uninformed player is not in a position to discover at the end of the game that his opponent has a different view of the game. The informed player may or

may not discover at the end of the game that he and his opponent have inconsistent beliefs about the game, depending on whether eq. (3-12) holds. Note that if both $p, q > 1/2$ or $p, q < 1/2$, the players never discover the differences in their models.

In this section, we presented and analyzed a simple class of two-person non-cooperative nonzero sum one stage subjective games of incomplete information, and showed how the inconsistent beliefs of the players lead to counterintuitive behavior. An important issue which has not been discussed so far is whether the differences in beliefs between the two players are amplified or smoothed out if the game is repeated over and over. We address this issue in the next section.

SECTION 4:

INFINITELY REPEATED SUBJECTIVE GAMES OF INCOMPLETE INFORMATION

In this section, we consider the infinitely repeated version of the class of games studied in section 3. First, we study the situation where no player is informed concerning the outcome of Nature's move. Then, we consider the case of private information. For this class of games, we show that repetition can alleviate differences between subjective models and lead to agreement about the outcome of the game. In addition, we show that the value of private information is always positive in this infinitely repeated game. Finally, we examine the effect of various bargaining models on the outcome of the game.

Before we proceed with the analysis, let's define the meaning of a solution to an infinitely-repeated nonzero sum two-person game. According to the results of [11],[16],[17], the set of equilibrium outcomes of the infinitely repeated game are all the payoffs which are individually rational [15]. Among these equilibrium outcomes, the set of efficient equilibria (also known as the core equilibria of the game) are the set of outcomes which are also Pareto optimal. A pair of equilibrium strategies in the infinitely-repeated game will be a solution if and only if it produces payoffs among the core equilibria of the infinitely-repeated game.

4.1 No Information

Assume that neither player has information concerning Nature's move. In this case, the one-stage subjective expected payoff matrix of each player can be computed as in the previous section. The set of obtainable payoffs for players 1 and 2 are given in figures 4-1 and 4-2 according to player 1's perception, and figures 4-3 and 4-4 according to player 2's perception. The individually rational outcomes for each player have been outlined in the shaded areas of the figures. It is easy to see that the set of equilibrium outcomes of the infinitely repeated game, as perceived by the two players, do not coincide. Moreover, the set of core equilibria are different according to the perception of the two players. As far as player 1 is concerned, the set of core equilibria contains only one possible outcome, $(pa + (1-p)b, pa + (1-p)b)$, achieved by the strategy (λ, σ) . The core of the game for player 2 consists of all the points of the lines AB and BC.

Let S^i denote the set of equilibrium strategy pairs, as perceived by player i , $i=1,2$. If either

player plays a strategy not in S^i , then the other player will detect an immediate inconsistency between the player's models and an equilibrium will not be reached. We can now define what are equilibrium strategies in this class of games.

Definition: The set S of equilibrium strategies of the infinitely- repeated subjective game of incomplete information is the intersection of S^1 and S^2 .

An immediate result follows from this definition:

Lemma 4.1 If a pair of strategies $(\gamma^1, \gamma^2) \in S$, the players never realize that they have different models.

Proof:

For any (γ^1, γ^2) in S , each γ^i [11] consists of two parts: The strategy player i implements as long as the other player does not deviate from his announced strategy, and the threat player i implements if the other player deviates from his announced strategy. We must show that, whether deviations from the announced strategy occur or not, no player finds out that they have different models.

Assume that both players follow their announced equilibrium strategies. Then, since $(\gamma^1, \gamma^2) \in S$, the resulting payoffs are individually rational according to each player's perception, so the players cannot detect that they have different models.

Assume that player 1 (2) deviates from his announced equilibrium strategy and player 2 (1) detects this deviation. Player 1's (2's) threat consists of selecting the strategy which reduces player 2's (1's) payoff the most; this strategy is $\mu(\tau)$. From player 2's (1's) perspective, player 1's (2's) threat strategy is also $\mu(\tau)$. Hence, when a deviation occurs the players cannot detect that they have different models, because their opponent's threat strategies are the same as those predicted by their own models. q.e.d.

4.2 Private information

Consider the infinite repetition of the game described in section 3, and assume that player 1 is privately informed of Nature's move. The main feature of this game is that the uninformed player can collect additional information throughout the play of the game by watching the behavior of the informed player. On the other hand, the informed player has to decide what information (if any) he has to reveal to his opponent and at what rate. Various interesting questions can be asked about this game:

1. Can revelation of private information be beneficial for the informed player?
2. What is the value of private information?
3. Does repetition alleviate the differences between the models of the two players?
4. Can the uninformed player take advantage of the private information of the informed player?

To find the answer to these questions, we shall determine core equilibrium strategies of the infinitely repeated subjective game. These core equilibrium strategies are strategies which result in core equilibrium payoffs according to both players' subjective model of the infinitely repeated game. The following results characterize these equilibrium payoffs.

Theorem 4.2: There is a unique pair of payoffs corresponding to core equilibrium strategies in the infinitely repeated subjective game when player 1 is privately informed of the outcome of Nature's move. This pair is $\{ra + (1-r)c, ra + (1-r)c\}$.

Before proving this theorem, let's describe how the core equilibrium payoffs are determined. First, the feasible region of payoffs for the game of incomplete information according to each player's perception is determined as the payoffs which are individually rational for each player [10]. Then, the core equilibrium payoffs of each feasible region is determined by identifying the Pareto optimal outcomes.

In order to develop the proof of theorem 4.2, we will need the following results concerning infinitely repeated nonzero sum games of incomplete information [10].

Lemma 4.3: Player 2 can limit the actual and perceived payoff of player 1 to c by playing the pure strategy τ throughout the repeated game. This strategy is the Blackwell strategy [18] of player 2.

Proof: See Appendix.

Lemma 4.4: According to player 1's perception, player 1 can limit the payoff of player 2 to $pc + (1-p)d$. The strategy achieving this payoff is the maxmin strategy for the following infinitely repeated zero sum game of incomplete information:

Game 1			Game 2		
	σ	τ		σ	τ
λ	$-a$	$-b$	λ	$-b$	$-a$
μ	$-c$	$-d$	μ	$-d$	$-c$

where $\text{Prob}(\text{Nature chooses game 1}) = p$.

Lemma 4.5: According to player 2's perception player 1 can limit the payoff of player 2 to $qd + (1-q)c$. The strategy achieving this payoff is the maxmin strategy for the infinitely repeated zero sum game of incomplete information of lemma 4.4, where

$\text{Prob}(\text{Nature chooses Game 1}) = q$.

Lemmas 4.4 and 4.5 are direct consequences of the definition of the threat strategy and the results of zero sum infinitely repeated games of incomplete information in [9].

Lemma 4.6: According to player 1's perception, the payoff $\{pa + (1-p)c, pa + (1-p)c\}$ is an individually-rational equilibrium payoff. According to player 2's perception, the payoff $\{qa + (1-q)c, qa + (1-q)c\}$ is an individually-rational equilibrium payoff.

Proof: Individual rationality follows from Lemmas 4.3 and 4.5. The above payoff is achieved by the following strategy: At the first stage of the game player 1 (the informed player) plays λ if Nature selects Game 1 and μ if Nature selects Game 2, thereby revealing Nature's move to player 2. Player 2 can play any strategy at stage 1. For all subsequent stages, if Nature chose Game 1, the players choose (λ, σ) . If Nature chose Game 2, the players choose (μ, τ) . If either player deviates, the other player implements his threat strategy of lemma 4.3 or 4.4. Note that neither

player 1 nor player 2 can improve his payoff by deviating from the announced strategy, since the threat strategy yields a lower payoff than the equilibrium for the deviating player; consequently, the strategies described above result in an equilibrium, according to each player's perception of the game. The values achieved by these strategies are the values postulated in the Lemma. q.e.d.

Lemma 4.7: According to player 1's perception, the payoff $p_a + (1-p)c$ is the most player 1 can achieve in the infinitely repeated game. According to player 2's perception, the payoff $q_a + (1-q)c$ is the most player 1 can achieve in the infinitely repeated game.

Proof: See appendix.

Proof of Theorem 4.2:

From Lemmas 4.6 and 4.7, the equilibrium payoffs of Lemma 4.6 are individually rational and Pareto optimal according to both players' perceptions. From lemma 4.7, these equilibrium payoffs are the unique payoffs in the core of the game according to each player's perception. In order to achieve these payoffs, player 1 must reveal his private information concerning Nature's move. Hence, all core equilibrium strategies for player 1 must reveal Nature's move in a finite number of repetitions. Once Nature's move is common knowledge, each player can use his dominant strategy in the appropriate Game, so that all core equilibrium strategies result in the unique equilibrium payoffs for the infinitely-repeated game. q.e.d.

Corollary 4.8: Assume both players follow the core equilibrium strategies described in the proof of lemma 4.6. Then, the players do not realize that they have different subjective models.

Proof: As long as neither player deviates from his equilibrium strategy, they cannot discover the differences in their model, since these strategies are core equilibrium strategies. Suppose player 1 deviates from his equilibrium strategy. Such a deviation is immediately detected by player 2, who switches to his threat strategy τ as described in lemma 4.3. This strategy is independent of the value of q , hence the players are unable to detect any difference in their subjective models. Suppose player 2 deviates from his equilibrium strategy. Such a deviation is immediately detected by player 1, who switches to his threat strategy as described in lemma 4.4. This strategy is described in [19] as follows: A set of lotteries is available to player 1. Depending on the outcome of Nature's move, a lottery is chosen by player 1. This lottery is performed and its outcome

determines the move of player 1 for the remainder of the game. In order for the two players to realize that they have different models, the move of player 1 must have zero probability according to the model of player 2. Using the results of [19], the maxmin strategy of player 1 is: According to player 1's perception, player 1 performs either a lottery whose outcome is λ with probability 1 or a lottery whose outcome is λ with probability 1/2. According to player 2's perception, player 1 performs either a lottery whose outcome is μ with probability 1 or a lottery whose outcome is λ with probability 1/2. Therefore, even when player 1 switches to his threat strategy, the players never realize that their subjective models are different because there is no contradictory behavior. q.e.d.

Corollary 4.9: For the game of theorem 4.2, the value of private information is positive for player 1.

Proof: The value of private information for player 1 can be computed explicitly using theorem 4.2. Since $pa + (1-p)b$ is the only core equilibrium of the infinitely-repeated game without private information for either player, the expected value of information to player 1 is the difference between the payoff of Theorem 4.2 and this payoff. This value is

$$V^1 = (1-p)(c-b) > 0.$$

Note also that

$$V^2 = (1-q)(c-b) > 0,$$

so the private information of player 1 also has positive value for player 2. q.e.d

In sum, we have answered many of the questions asked at the beginning of this subsection. Specifically, repetition of the game serves to alleviate the differences in the models of the players and leads to eventual agreement on a common pair of strategies, and the value of private information is positive for the informed player, unlike the results for the static game of the previous section. Note that, due to the symmetry in the game, a similar set of lemmas and theorems can be established if player 2 were the informed player. When one player has private information, the perceived core of the game for each player has a single pair of equilibrium payoffs, and repetition of the game serves to eliminate the differences in initial probability assessments, in contrast with the

situation when neither player has private information, where the perceived core of the game for player 2 can contain many additional equilibrium payoffs. The selection of equilibrium strategies in this case is the result of bargaining between the two players to decide which pair of core equilibrium payoffs they will achieve. In the next subsection, we examine the effect of various bargaining theories on the outcome of the infinitely repeated subjective game without private information.

4.3 The effect of Bargaining Theories on the Outcome of a Subjective Infinitely Repeated Game of Incomplete Information

In this section, we investigate the effect of the bargaining theories of Nash [21] and Zeuthen [22,23] on the outcome of the infinitely repeated game of subsection 4.1. The core equilibrium payoffs of the game according to player 1's perception are the single point $(pa + (1-p)b, pa + (1-p)b)$. According to player 2's perception, the core equilibrium payoffs of the game are the pairs of points on the boundary of the shaded region of fig. 4-2. We show that, depending on the bargaining model used by the players, they may or may not realize the difference in their models during the course of bargaining.

Theorem 4.10: Consider the bargaining problem for the infinitely repeated game of Section 4.1. If the Nash bargaining model with moves consisting of the choice of strategies is used, the players will agree on the core equilibrium strategy and never realize the differences in their models.

Proof: Nash's bargaining model is a game in normal form where each player has only one move. In this move, the player announces a strategy which achieves the payoffs corresponding to the Nash bargaining solution, satisfying Nash's axioms [21]. The maxmin values which each player can guarantee himself are $(pd + (1-p)c, pd + (1-p)c)$ according to player 1's perception, and $(qc + (1-q)d, qc + (1-q)d)$ according to player 2's perception. In either case, the maxmin values lie on the line $x=y$, as do points B and B' in figures 4-1 and 4-2. Hence, the payoffs corresponding to Nash's bargaining solution correspond to points B and B' in these figures. The strategies which achieve these payoffs are the pair of strategies (λ, σ) , according to each player's perception. Thus, the players will agree on the strategies announced and never realize the difference in their underlying perceptions of the game.

In the above result, players only exchanged their final bargaining strategies. When utility is linearly transferable between the players, Nash's bargaining model with threats can be used. In this model, players exchange threat strategies, then select the bargaining solution based on these threat strategies. The players can detect the inconsistency between their models if the threat strategies announced are inconsistent with the players' models.

Theorem 4.11: Assume that players use Nash's bargaining model with threats and transferable utilities[28]. If

$$p \leq \frac{(a-d+b-c)}{a-d} = p^*$$

then the players do not detect that they have different perceptions of the game. If $p > p^*$, the players detect an inconsistency in their perceptions of the game.

Proof: In Nash's bargaining model with threats and transferable utilities[28], the threat strategies for the players, according to player 1's perception, are the solutions of the zero sum game

	σ	τ
λ	0	$p(c-b) + (1-p)(d-a)$
μ	$-p(c-b) + (1-p)(d-a)$	0

If $p < p^*$, the optimal pair of threat strategies is (μ, τ) . If $p > p^*$, the optimal pair of threat strategies is (λ, σ) . According to player 2, since $q < 1/2$, the optimal pair of threat strategies is (μ, τ) . Hence, if $p > p^*$, the players will detect an inconsistency in the announced threat strategies. If $p < p^*$, the threat strategies announced by each player will be consistent, and result in expected payoffs which are on the $x=y$ line. Hence, the Nash bargaining solution will be the same as the previous theorem, and the players will not detect the difference in their perceptions of the game.

Other bargaining models can lead to players discovering the inconsistencies in their models. Consider Zeuthen's regular bargaining model [22]-[23], which is a game in extensive form which

allows bargaining to proceed in steps. We assume again that the moves of players are strategies corresponding to the players' payoff demands. Then, if the players do not reach an agreement in one move, it is possible that player 2 (whose perception of the game is described by figure 4-2) may propose a move which player 1 cannot interpret within the terms of his own model (e.g. player 2 may propose a move corresponding to a payoff which is in the core of his own game but not in the set of equilibria of the game perceived by player 1). At that point, player 1 has a different model of the game since, by assumption, he excludes the possibility of an irrational opponent.

Theorems 4.10 and 4.11 illustrate the role of the bargaining model on the outcome of the subjective game. The outcome of the infinitely repeated subjective game depends on the number of steps required to reach an agreement. If the model predicts that agreement is reached in one step, each player may not have the opportunity to realize that the intended meaning of the move of the other player was different from what he perceived it to be. On the other hand, if the model predicts that more than one step may be required to reach an agreement, then the players may have the opportunity to realize that their subjective perceptions of the game are different.

SECTION 5. CONCLUSIONS

In this report, we formulated a class of "subjective games," where the players have different perceptions of the rules of the game and are unaware of the differences in their perceptions. This class of games is a generalization of the team problem of asymptotic agreement studied in [14]. We developed a conceptual and analytical framework for studying the effects of these differences in perception on the strategies used by the players. By studying in detail a specific class of symmetric games of incomplete information, we showed that the properties of these "subjective games" are different from the properties of similar Bayesian games. Specifically, many features of the Bayesian games, such as the positive value of private information in symmetric games, are not maintained when the players' perceptions of the game are allowed to differ.

An important issue which arose from our formulation was whether the players discover that their perceptions are different during the play of the game. We showed that, in a static game, players often discover at the end of the game that they have different perceptions. Infinite repetition of the game, however, may alleviate the differences in the players' models, and lead to strategies where the players do not discover that they had different initial perceptions. In addition, we showed that, in an infinitely-repeated game, agreement on a core equilibrium strategy depends on the bargaining model adopted by the players in the game.

The rudimentary investigation reported here needs to be carried further. It is important to characterize the classes of subjective games where the rational strategies are insensitive to the differences in the perceptions of the players. In addition, our analysis should be extended beyond the point where the two players reach an impasse. To proceed further, it may be appropriate to formulate the conflict situation as a bargaining problem which is perceived differently by each player. For such a bargaining problem, the players seek strategies which belong to the core of all games. The bargaining theories of Harsanyi [23], Hearn [25], Kalai-Owen-Maschler [24], Owen [26] and Zeuthen [22] may prove useful in determining such strategies and could be the starting point for further investigation.

APPENDIX

Proof of Lemma 4.3:

To determine the payoff to which player 1 can be limited by player 2, it suffices to determine the approachable set for player 2 when he tries to minimize player 1's payoff. We define the following payoff matrix

$$G = \begin{array}{cc} & \begin{array}{c} [a,b] \\ [b,a] \end{array} & \begin{array}{c} [c,d] \\ [d,c] \end{array} \\ \begin{array}{c} [a,b] \\ [b,a] \end{array} & & \end{array}$$

The first (second) component of each entry of G gives the payoff to player 1 when game 1 (2) is played. These payoffs are shown in figure A-1. The approachable set for player 2 is the shaded area of payoffs of fig. A-1 [9]. A point in this area is guaranteed by using the pure strategy τ , which is the Blackwell strategy [9] for player 2. When player 2 uses this strategy, player 1's payoff cannot exceed c . q.e.d.

Proof of Lemma 4.7:

The maximum payoff which player 1 can hope to achieve is an equilibrium payoff which is individually rational for player 2. To determine this payoff, we define the payoff matrix

$$\begin{array}{cc} [sa+(1-s)b, sa+(1-s)b] & [sc+(1-s)d, sb+(1-s)a] \\ [sb+(1-s)a, sc+(1-s)d] & [sd+(1-s)c, sd+(1-s)c] \end{array}$$

and let s vary from 0 to 1. For

$$0 \leq s \leq \frac{(c-b)}{(c-b+a-d)} \quad (A-1)$$

we obtain

$$sb+(1-s)a \geq sd+(1-s)c \geq sa+(1-s)b \geq sc+(1-s)d. \quad (A-2)$$

From Lemma 4.5, the individually-rational payoff for the game with probability s is $sd + (1-s)c$. As illustrated in figure A-2, the maximum payoff for player 1 when A-1 holds is $(1-s)c + sd$.

For s in the range

$$\frac{(c-b)}{(c-b+a-d)} \leq s \leq 1/2 \quad (A-3)$$

we get

$$sb+(1-s)a \geq sa+(1-s)b \geq sd+(1-s)c \geq sc+(1-s)d. \quad (A-4)$$

Figure A-3 illustrates the maximum payoff which can be obtained by player 1. After some algebra, this payoff can be determined as

$$\frac{s^2(a-b)(3d-3c+b-a) + s\{(a-b)(a+3c-b-3d) - b(b+c) + ad + b^2\} + ab+bc-ac-bd}{s(a-b+d-c) + b - d} \quad (A-5)$$

For $s > 1/2$, the situation is illustrated in fig. A-4. In this case, the maximum payoff which player 1 can achieve is $sa + (1-s)b$.

According to [9], [20], the maximum payoff player 1 can achieve as s increases from 0 to 1 is given by the least concave function which majorizes the payoff function as s varies from 0 to 1. As illustrated in fig. A-5, this function is described by

$$sa + (1-s)c. \quad (A-6)$$

Consequently, the maximum payoff player 1 can achieve when $s = p$, according to his perception, is given by $pa + (1-p)c$. The maximum payoff he can achieve according to player 2's perception is given by $qa + (1-q)c$. q.e.d.

REFERENCES

1. Aumann, R.J. and M. Maschler, "Some Thoughts on the Minimax Principle," Management Science, Vol. 18, No. 5, Jan. 1972.
2. Ho, Y.C., "On the Minimax Principle and Zero-Sum Stochastic Differential Games," Tech. Report No. 630, Harvard University, Cambridge, MA, April 1972.
3. Kreps, D.M. and R. Wilson, "Sequential Equilibria," Stanford University Report, Dec. 1980.
4. Wilson, R., "Perfect Equilibria and Sequential Rationality," presented at Luminy Conference on Games with Incomplete Information, Sept. 1981.
5. Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," Int. Journal of Game Theory, Vol. 4, no.1.
6. Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," Zeit. für die Gesamte Staatswissenschaft, Vol. 12.
7. Harsanyi, J.C., "Games of Incomplete Information Played by Bayesian Players," Parts I-III, Management Science, Vol. 14, nos. 3,5,7, 1967-1968.
8. Aumann, R.J., M. Maschler et. al., "Game Theoretic Aspects of Gradual Disarmament," In: Mathematica report to Arms Control & Disarmament Agency, 1966-1968.
9. Sorin, S. "An Introduction to Two-Person Zero-Sum Repeated Games of Incomplete Information," Technical Report no. 312, I.M.S.S.S., Stanford University, May 1980.
10. Hart, S., "Non-Zero Sum Two-Person Repeated Games with Incomplete Information," Technical Report no. 367, I.M.S.S.S., Stanford University, Feb. 1982.
11. Aumann, R.J., "Acceptable Points in General Cooperative n-Person Games," Contributions to the Theory of Games, Vol. IV, Annals of Mathematical Studies, No. 40, Princeton: Princeton Univ. Press.
12. Von Neumann, J. and Morgenstern, O., Theory of Games and Economic Behavior, J. Wiley, 1944.
13. Savage, L., The Foundations of Statistics, Dover 1972.
14. Teneketzis, D. and P. Varaiya, "Consensus in Distributed Estimation with Inconsistent Beliefs," Systems and Control Letters, July 1984.

15. Levine, P., and J.P. Ponssard, "The Values of Information in Some Nonzero Sum Games," International Journal of Game Theory, Vol. 6, No. 4.
16. Hart, S., Lect. Notes: Special Topics in Game Theory, Stanford University, Stanford, CA 1979.
17. Aumann, R., Lectures On Game Theory, Stanford University, Stanford, CA, 1976.
18. Blackwell, D., "An Analog of the Minimax Theorem for Vector Payoffs," Pacific J. of Math., Vol. 6, 1956.
19. Ponssard, J.P. and S. Zamir, "Zero-Sum Sequential Games with Incomplete Information," Int. Journal of Game Theory, Vol. 2.
20. Zamir, S., "Convexity and Repeated Games," in J.P. Aubin (Ed.) Analyse Convexe et ses Applications, Lecture Notes in Economics and Math. Systems, No. 102, Springer-Verlag.
21. Nash, J.F., "The Bargaining Problem," Econometrica, vol. 18, 1950.
22. Zeuthen, F., Problems of Monopoly and Economic Warfare, Routledge and Kegan Paul: London, 1930.
23. Harsanyi, J.C., Rational Behavior and Bargaining Equilibrium in Games and Social Situations, Cambridge Univ. Press, 1977.
24. Kalai, G., M. Maschler and G. Owen, "Asymptotic Stability and Other Properties of Trajectories and Transfer Sequences Leading to the Bargaining Sets," Tech. Report 93, Economics Series, I.M.S.S.S., Stanford University, 1973.
25. Stearns, R.E., "Convergent Transfer Schemes for n-Person Games," Trans. American Math. Soc., Vol 134, 1968.
26. Owen, G., "A Note on the Convergence of Transfer Schemes in n-person Games," Int. Journal of Game Theory, vol. 4.
27. Tversky, A. and D. Kahneman, "Judgement under Uncertainty: Heuristics and Biases," Science, Vol. 185, 1974.
28. Owen, G., Game Theory, W.B. Saunders, Philadelphia, 1968.




Figure 4-1. Individually Rational Payoffs According to Player 1, $p > p^*$.




Figure 4-2. Individually Rational Payoffs According to Player 1, $p < p^*$.




Figure 4-3. Individually Rational Payoffs According to Player 2, $q < 1 - p^*$.




Figure 4-4. Individually Rational Payoffs According to Player 2, $q > 1 - p^*$.




Figure A-1: Payoff to Which Player 1 can be Limited by Player 2




Figure A-2. Maximum Payoff of Player 1 when $0 \leq s \leq 1 - p^*$.




Figure A-3. Maximum Payoff of Player 1 when $1 - p^* \leq s \leq 1/2$.




Figure A-4. Maximum Payoff of Player 1 when $1/2 \leq s \leq 1$.




Figure A-5. Maximum Payoff Achievable by Player 1.

REFERENCES

1. Von Neumann, J. and O. Morgenstern, Theory of Games and Economic Behavior, J. Wiley, 1944.
2. Levine, P., and J.P. Ponssard, "The Values of Information in Some Nonzero Sum Games," International Journal of Game Theory, Vol. 6, No. 4.
3. Savage, L., The Foundations of Statistics, Dover, 1972.
4. Einhorn, H.J. and R.M. Hogarth, "Behavioral Decision Theory: Processes of Judgement and Choice," Ann. Review Psychology, 1981.
5. Slovic P., B. Fischhoff, and S. Lichtenstein, "Behavioral Decision Theory," Ann. Review Psychology, Vol. 28, 1977.
6. Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory, Vol. 4, No. 1, pp. 25-55, 1975.
7. Myerson, R.B., "Refinements of the Nash Equilibrium Concept," Intern. Journal of Game Theory, Vol. 7, No. 2, pp. 73-80, 1978.
8. Okada, A., "On Stability of Perfect Equilibrium Points," Intern. Journal of Game Theory, Vol. 10, No. 2, pp. 67-73, 1981.
9. Aumann, R.J. and M. Maschler, "Some Thoughts on the Minimax Principle," Management Science, Vol. 18, No. 5, January 1972.
10. Ho, Y.C., "On the Minimax Principle and Zero-Sum Stochastic Differential Games," Technical Report No. 630, Harvard University, Cambridge, MA., April 1972.
11. Castanon, D.A., "Equilibria in Stochastic Dynamic Games of Stackelberg Type," Ph.D. Thesis, Dept. of Mathematics, MIT, 1976.
12. Harsanyi, J.C., "Games of Incomplete Information Played by Bayesian Players," Parts I-III, Management Science, Vol. 14, No. 3, No. 5, No. 7, 1967-1968.
13. Harsanyi, J.C., Rational Behavior and Bargaining Equilibrium in Games and Social Situations, Cambridge University Press, 1977 (Chapter 8).

14. Ponssard, J.P. and S. Zamir, "Zero-Sum Sequential Games with Incomplete Incomplete Information," International Journal of Game Theory, Vol. 2, pp. 99-107.
15. Sorin, S., "An Introduction to Two-Person Zero-Sum Repeated Games of Incomplete Information," Technical Report No. 312, I.M.S.S.S., Stanford University, May 1980.
16. Hart, S., "Non-Zero-Sum Two-Person Repeated Games of Incomplete I Information," Technical Report No. 367, Institute for Mathematical Studies in the Social Sciences, Stanford University, February 1982.
17. Kadane, J.B. and P.D. Larkey, "Subjective Probability and the Theory of Games," Management Sci., Vol. 28, No. 2 (1982), pp. 113-120.
18. "The Confusion of Is and Ought in Game Theoretic Contexts," Management Sci., Vol. 29, No. 12, (1983), pp. 1365-1379.
19. Wilson, J.G., "Subjective Probability and the Prisoner's Dilemma," Management Science, January 1986.
20. Wohl, J.G., "Force Management Decision Requirements for Air Force Tactical Command and Control," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-11, 1981.
21. Tversky, A. and D. Kahneman, "Prospect Theory: An Analysis of Decision and Risk," Econometrica, 1979, Vol. 47.
22. Luce, R.D., "The Choice Axioms after 20 Years," J. Math. Psych., Vol. 15, 1977.
23. Luce, R.D. and H. Raiffa, Games and Decisions, J. Wiley, 1957.
24. Machina, M., "The Economic Theory of Individual Behavior Toward Risk: Theory, Evidence and New Directions," IMSSS Report 433, Stanford University, October 1983.
25. Harsanyi, J.C. "Subjective Probability and the Theory of Games: Comments on Kadane and Larkey's Paper," Management Science, Vol. 28, 1982.
26. Kadane, J.P., "On Choosing Between Rev. Bayes and Prof. Von Neumann," Management Science, Vol. 29, 1983.
27. Roth, A.E. and F. Schoumaker, "Subjective Probability and the Theory of Games: Some Further Comments," Management Science, Vol. 29, 1983.
28. Roth, A., M.H., "Modeling Semirational Cognitive Behavior," Management Science, Vol. 29, 1983.
29. Aumann, R.J., "Agreeing to Disagree," Annals of Statistics, Vol. 4, No. 6, 1976, pp. 1236-1239.

30. Borkar, V. and P. Varaiya, "Asymptotic Agreement in Distributed Estimation," IEEE Transactions on Automatic Control, Vol. AC-27, No. 3, 1982.
31. Tsitsiklis, J.N. and M. Athans, "Convergence and Asymptotic Agreement in Distributed Decision Problems," IEEE Transactions on Automatic Control, Vol. AC-29, No. 1, 1984, pp. 42-50.
32. Geanakoplos, J.D. and H.M. Polemarchakis, "We Can't Disagree Forever," Journal of Economic Theory, Vol. 26, 1982, pp. 192-200.
33. Washburn, R.B. and D. Teneketzis, "Asymptotic Agreement Among Communicating Decisionmakers," Stochastics, Vol. 13, 1984, pp. 103-129.
34. Teneketzis, D. and P. Varaiya, "Consensus in Distributed Estimation With Inconsistent Beliefs," Systems and Control Letters, Vol. 4, 1984, pp. 217-221.
35. Athans, M., "The Expert Team of Experts Approach to C^2 Organization," Proceedings of the 4th MIT-ONR Workshop on C^3 Problems, Volume 4, October 1981.